Journal of Artificial Intelligence Research 48 (2013) 175-230 Submitted 05/13; published 10/13

An Online Mechanism for Multi-Unit Demand and its Application to Plug-in Hybrid Electric Vehicle Charging

Valentin Robu VR2@ECS.SOTON.AC.UK
Enrico H. Gerding EG@ECS.SOTON.AC.UK
Sebastian Stein SS2@ECS.SOTON.AC.UK
University of Southampton, SO17 1BJ, Southampton, UK

David C. Parkes PARKES@EECS.HARVARD.EDU
Harvard University, Cambridge, MA 02138, USA

Alex Rogers ACR@ECS.SOTON.AC.UK
University of Southampton, SO17 1BJ, Southampton, UK

Nicholas R. Jennings NRJ@ECS.SOTON.AC.UK
University of Southampton, SO17 1BJ, Southampton, UK
King Abdulaziz University, Jeddah, Saudi Arabia

Abstract

We develop an online mechanism for the allocation of an expiring resource to a dynamic agent population. Each agent has a non-increasing marginal valuation function for the resource, and an upper limit on the number of units that can be allocated in any period. We propose two versions on a truthful allocation mechanism. Each modifies the decisions of a greedy online assignment algorithm by sometimes cancelling an allocation of resources. One version makes this modification immediately upon an allocation decision while a second waits until the point at which an agent departs the market. Adopting a prior-free framework, we show that the second approach has better worst-case allocative efficiency and is more scalable. On the other hand, the first approach (with immediate cancellation) may be easier in practice because it does not need to reclaim units previously allocated. We consider an application to recharging plug-in hybrid electric vehicles (PHEVs). Using data from a real-world trial of PHEVs in the UK, we demonstrate higher system performance than a fixed price system, performance comparable with a standard, but non-truthful scheduling heuristic, and the ability to support 50% more vehicles at the same fuel cost than a simple randomized policy.

1. Introduction

Designing mechanisms for allocating scarce resources to self-interested agents is a central research topic in artificial intelligence (Sandholm, 2002; Engel & Wellman, 2010). The aim of this work is to devise mechanisms that satisfy certain desirable properties, such as truthfulness and efficiency. Many settings where such mechanisms can be applied are characterised by dynamic supply and demand, i.e., agents arrive and leave the market over time, and the availability of supply also changes over time. This has led the field of *online mechanism design*, in which agents are incentivised to report truthfully not only their value for a given allocation, but also the period they are available in the market (Parkes, 2007). However, to date, most of the existing work in this field assumes that the valuations of the agents for a certain allocation can be described by a single parameter, the so-called *single-valued* domains. Existing approaches that do consider multi-valued domains rely on access to a probabilistic model of supply and demand and the ability to compute an optimal allocation policy, which becomes computationally infeasible for realistic settings.

To address these shortcomings, we extend the state of the art by developing a novel model-free mechanism (i.e., which assumes no knowledge of future demand or supply) for multi-valued demand. In particular, we consider domains with multi-unit demand and agents with *non-increasing marginal values*. In such domains, the first units allocated to an agent have a higher (or equal) marginal value for this agent compared to any subsequent units. In the online settings we consider, resources are continuously produced and *perishable*, thus the available supply must be allocated in each period. Moreover, the supply available in each period is not known in advance, but only at the start of that period.

Examples of settings with non-increasing marginal values and perishable resources occur in many real-life settings. One such example is cloud computing, where jobs arrive over time and perishable computational resources must be allocated to these jobs (Porter, 2004; Stein, Gerding, Rogers, Larson, & Jennings, 2011). In particular, the non-increasing marginal value model applies naturally to large-scale data processing or optimisation with any-time computation. In such a setting, the first unit of computation provides a solution of a certain quality, while subsequent units allow improving this solution, up to a level when further computation is no longer useful. Hence, the first units are more valuable, as they already provide a good approximation of the desired solution, while subsequent units increase this value, but by a marginally non-increasing amount. Another example is online advertising, where impressions need to be allocated as soon as users visit a webpage (Constantin, Feldman, Muthukrishnan, & Pal, 2009). The non-increasing marginal values also applies in this setting, since additional exposure of a set of users to the same ad will likely have a decreasing impact.

A third example, studied extensively in this paper, is the allocation of electricity for the charging of plug-in hybrid electric vehicles (PHEVs). Similar to pure (non-hybrid) electric vehicles (pure EVs), these vehicles can be charged directly at an electric charging point. The difference is that PHEVs have both an electric motor and an internal combustion engine, and are widely seen as a solution to the problem of *range anxiety*, i.e., fear that a car will run out of electricity in the middle of nowhere (Eberle & von Helmolt, 2010).[1] However, with the associated increase in demand for electricity, there are significant concerns within the electricity distribution industries regarding the widespread use of such vehicles, since the high charging rates that PHEVs require (up to three times the maximum current demand of a typical home) could overload local electricity distribution networks at peak times (Fairley, 2010). One approach to address this concern (e.g., adopted by the Pacific Gas and Electric Company in California) is to introduce time-of-use pricing plans that seek to shift demand. A more sophisticated approach that takes into account the valuations of self-interested owners, is to design an online mechanism which schedules access dynamically in order to prevent network overload. The assumption of decreasing marginal values is justified here

1. Practical examples of PHEVs include versions of cars such as the Toyota Prius and Honda Insight, that can drive on petrol, but whose batteries can also be charged directly at an electrical charging point.

because a vehicle owner is more likely to use the first units of electricity, and can always use the combustion engine as an alternative in case she runs out of electricity (and so the car can still be used even if it is not fully charged). Against this background, the main contributions of our work are:

- We develop a new model-free online mechanism for settings where participating agents have non-increasing marginal values for units of a perishable good. We adopt a greedy algorithm coupled with a method to modify the allocation to ensure incentive compatibility. This involves cancelling part of the proposed allocation, and we explore two ways of performing this cancellation: *immediately*, i.e., at each time step before a resource is actually allocated, or *on departure* of an agent from the market (at which point we must take back allocated items). Both variants are (weakly) *dominant-strategy incentive compatible* (DSIC), so that participants have no incentive to misreport their valuations or arrival-departure dynamics.

- We analyse the worst-case performance achieved by the mechanisms relative to the optimal offline allocation, considering both the number of units that need to remain unallocated in order to achieve incentive compatibility and the total value of the allocation (i.e., the allocative efficiency). The variation with on-departure cancellation results in higher allocative efficiency, and is more tractable, but may involve additional practical challenges. For example, in the PHEV charging domain, this occasionally requires a vehicle's battery to be partially discharged prior to departure.

- We evaluate the online mechanism through numerical simulations of an abstract domain and the PHEV charging domain, and compare the results to several benchmarks that assume non-strategic agents, including an optimal offline solution, a scheduling heuristic and a greedy algorithm without cancellation. The simulations of the PHEV domain are based on real data from a large-scale trial of (pure) EVs in the UK. Valuations are derived from real monetary savings, by considering factors such as fuel prices, the distance that the owner expects to travel, and the energy efficiency of the vehicle. The results establish that the mechanism outperforms a fixed-price mechanism in terms of allocative efficiency in both domains, while performing only slightly worse than the non-incentive compatible scheduling solutions. In addition, the mechanism with on-departure cancellation scales easily to hundreds of agents.

We focus on allocative efficiency rather than revenue, as this is appropriate to many domains of interest. For example, in the PHEV charging domain, it is reasonable that the goal is to allocate capacity efficiently in order to maximise value to the user base of a power company, given the significant constraints on charging capacity. Moreover, in many cloud computing applications (for example, in large-scale scientific computing), the goal is to allocate capacity to the jobs that are most urgent or important. However, if in practice the seller wants to guarantee a a minimum revenue from each unit sold, it would be easy to include a reserve price. If this minimum price is set the same for all units and all time points, this would not affect the properties of our mechanism.

The remainder of this paper is organised as follows. We first discuss related work (Section 2), before formally introducing the model (Section 3). In Section 4, we define the online mechanisms and study their strategic properties. Then, in Section 5, we develop the worst-case analytical results, followed in Section 6 by a discussion of how to compute the allocations and payments in practice. Using both real and synthetic input data, we present the results of the experimental evaluation of our mechanisms in Section 7, and we conclude in Section 8.

2. Related Work

In this section, we first review existing work on online mechanism design (Section 2.1), and then provide a background to the PHEV charging application, along with an overview of previous work that considers this problem (Section 2.2).

2.1 Online Mechanism Design

One line of work in online mechanism design aims to develop online variants of the Vickrey-Clarke-Groves (VCG) mechanism. In this context, Parkes and Singh (2003) consider the problem of maximising the long-term allocative efficiency of a system of self-interested agents that arrive and depart dynamically. They model the online mechanism design problem as a Markov decision process (MDP), whose solutions can be used to implement optimal policies in a truth-revealing Bayes-Nash equilibrium. In related work, Gershkov and Moldovanu (2010) examine the allocation of a set of goods to a dynamic population of randomly arriving buyers. They consider two settings: one in which there is a common deadline for allocating objects to all buyers, and a second one without a firm deadline, but in which buyers are impatient, assigning higher value to items allocated sooner.

Unlike both Parkes and Singh (2003), and Gershkov and Moldovanu (2010), the mechanism proposed in this paper is "model-free" (which has the advantage that no prior knowledge or distribution is required about the other agents' types or future allocations), and we focus on the stronger concept of dominant-strategy incentive compatibility (where reporting truthfully is a best response regardless of what other agents are doing, even if they are irrational). Such an approach requires fewer assumptions, and makes computing allocations more tractable compared to VCG-like approaches. This is because VCG generally requires the allocations to be optimal in expectation (perhaps in a constrained space of policies), whereas, as we will show, we can use greedy heuristics.

Model-free settings are considered by Hajiaghayi et al. (2005), Parkes (2007) and Porter (2004). The work of Porter examines the scheduling of jobs on a single machine and proposes an incentive compatible mechanism for this setting. However, his work assumes a setting where the results of a job are released to an agent only on completion or by the agent's reported deadline. While this assumption is reasonable for scheduling computational jobs on a server, it is not suitable for our setting, since the "goods" (i.e. electricity units) must be allocated instantly when they become available.

The work which considers an online setting most similar to the one we consider is by Hajiaghayi et al. (2005). They study the problem of online scheduling of a single, re-usable resource over a finite time period, in which each agent has an arrival-departure window when they are active in the market. Agents may misreport both their valuation, as well as their arrival and departure, subject to an assumption of "limited misreports" (i.e., no early arrival or later departure misreports are possible). For this setting, they characterise truthful allocation and payment policies, and prove worst-case approximation ratios with respect to the optimal offline allocation. A key limitation is that the mechanism proposed by Hajiaghayi et al. concerns single-valued domains, whereas we consider a multi-unit setting with decreasing marginal values. We show that their mechanism does not directly apply to the multi-unit case, requiring, in some cases, additional cancellation rules to be applied to ensure truthfulness.

Multi-unit demand is considered in the work of Lavi and Nisan (2004), who propose an online auction model in which the mechanism is required to make decisions about each bid as it is received. They provide a characterisation of incentive compatibility in such domains in terms of supply curves, a concept which relates closely to our threshold mechanism characterisation. However, in their online auctions model, the auctioneer must respond to each bid immediately, before considering other bids. In this response, the mechanism determines both the quantity to be sold and the price to be paid. This would not be applicable to the setting presented in this paper, where there is a limit on the number of perishable units that can be allocated in each time interval. Moreover, the window-based allocation allows the prices to be determined dynamically, based on the bids observed from the other agents until departure. In a similar vein, Babaioff, Blumrosen, and Roth (2010) consider an online auction model where future supply is unknown, and characterise several subclasses of truthful mechanisms. Their domain is different from ours, as bidders in their model do not specify multi-dimensional demands and non-increasing marginal values.

Other related work on online mechanism design adapts the *consensus* algorithm for online stochastic optimisation proposed by Bent and Van Hentenryck (2004) to a setting with self-interested agents. In this context, Parkes and Duong (2007), and Constantin and Parkes (2009) first propose the idea of modifying the decision of an algorithm by cancelling part of the allocation in order to ensure incentive compatibility. Unlike the present paper, the setting assumes single-valued private information and the approaches are not applicable to agents with non-increasing marginal values. Also for single-valued settings and for the pure EV domain, Stein, Gerding, Robu, and Jennings (2012) propose a model-based online mechanism that assumes knowledge of future supply and uses pre-commitment to ensure online allocations are truthful.

2.2 Electric Vehicle Charging

Multi-agent systems and AI techniques are increasingly used to address challenges in the Smart Grid (e.g., Vytelingum, Voice, Ramchurn, Rogers, & Jennings, 2011; Robu, Kota, Chalkiadakis, Rogers, & Jennings, 2012), and EV charging is one of the most important application areas. Work on the automatic scheduling of EV charging typically allows individual vehicle owners to indicate the times at which the car will be available for charging, enabling automatic scheduling while satisfying the constraints of the distribution network. In this vein, Clement, Haesen, and Driesen (2009) propose a centralised scheduler, which makes optimal use of the network capacity when vehicle owners report their expected future vehicle use to the system. Sundstrom and Binding (2012) tackle the problem of charging multiple electric vehicles considering distribution grid constraints, formalise the underlying optimisation problem and propose a novel method based on load flow to solve it. But strategic behaviour remains possible in these approaches; e.g., an owner may indicate an earlier departure time or further travel distances in order to receive preferential charging. This can result in a high cognitive load for car owners, and may lead to inefficient schedules that are not based on actual user requirements, leading to an efficiency loss.

The potential for speculation by strategic agents has been identified as a crucial problem in other scheduling problems, such as scheduling computational jobs on a cluster (Porter, 2004), scheduling of computation-intensive services on the cloud (Stein et al., 2011) or market-based scheduling of loads in transportation logistics (Robu, Noot, La Poutré, & van Schijndel, 2011). With the increase in the number of EVs requiring charging, the potential for manipulation will become an increasingly pressing problem in PHEV scheduling as well.

Other approaches to EV scheduling include the lottery-based solution proposed by Vasirani and Ossowski (2011), in which the decision of whether to charge a vehicle or not is determined through a lottery system, designed to ensure a level of fairness in the resulting allocation. Unlike our work, however, the authors do not use game-theoretic principles to prove that participating vehicles have an incentive to report their preferences truthfully, thus in their scheme agents may have an incentive to speculate. Moreover, the experimental analysis reported by Vasirani and Ossowski does not adopt real data to derive EV driving patterns, charging capacities or network constraints.

Other recent work investigates using grid-integrated electric vehicles (GIVs) to sell power and storage capacity back to the grid – a concept known as vehicle-to-grid (V2G) (c.f. Kamboj, Kempton, & Decker, 2011). That work is different from ours, as it does not study the problem of coordinated charging of PHEVs under local network capacity constraints. In subsequent work to the model in the present paper (which first appeared as Gerding, Robu, Stein, Parkes, Rogers, & Jennings, 2011, and Robu, Stein, Gerding, Parkes, Rogers, & Jennings, 2011), we study the problem of charging *pure* EVs, that must receive a set amount of charge, otherwise they derive no value from the allocation (Stein et al., 2012). A second paper studies the problem of two-sided markets, where both PHEVs and charging stations compete to be matched (Gerding, Stein, Robu, Zhao, & Jennings, 2013). Unlike the present model, these papers only consider single-minded bidders, and the work of Stein et al. assumes access to a probabilistic model of the environment.

3. The Model

We consider an online mechanism design setting with discrete time steps, where in each period, multiple indivisible units of a perishable good are being sold, and each agent requires multiple units within a certain period. As we will show in Section 3.1, this model can also be used for continuously available resources, such as electricity and computational resources. In that case, an allocation decision consists of the amount of resource to be consumed by each agent over the next period until the following time point.

For convenience, an overview of the notation is provided in Table 1. Formally, let $S(t)$ denote the supply available at time t. Let $I(t) = \{1, 2, \ldots\}$ denote the set of agents that are in the market at time t or have already left the market. We do not assume access to a probabilistic model of future arrivals, departures or future supply beyond the current time period t. Agents are numbered according to their arrival time. An agent $i \in I(t)$'s *type* is described by the tuple $\theta_i = \langle \mathbf{v}_i, a_i, d_i, r_i \rangle \in \Theta$, where $\mathbf{v}_i$ is the *marginal valuation* vector, a_i and d_i, with $d_i \geq a_i$, are the *arrival* and *departure* times (the earliest and latest times that the agent is available in the market), r_i is the *maximum consumption rate* (i.e., the maximum number of units agent i can consume at any time t), and Θ is the set of all *admissible* types. Upon arrival, an agent needs to report a valuation function and a maximum consumption rate. These two aspects of an agent's reported type are required to remain unchanged while the agent is present, although her departure time can be modified (it only becomes known to the mechanism on the actual departure).

Each element $v_{i,k}$ of valuation $\mathbf{v}_i$ is called a *marginal valuation*, and represents the agent's *willingness to pay* for the k^{th} unit of the good, given that it has acquired $k-1$ units. We require:

Assumption 1. *Marginal valuations are non-increasing, i.e.,* $\forall i, k : v_{i,k} \geq v_{i,k+1}$.

Given this, an agent's utility function is $U(k, x, \theta_i) = V(k, \theta_i) - x$, where x is the agent's payment to the mechanism, $V(k, \theta_i) = \sum_{j=1}^{k} v_{i,j}$ is the total value derived given its type, and k is the number of units allocated to this agent between its arrival and departure.

A *mechanism* asks agents to report their types and, based on this information, decides on an allocation of the supply and a payment for units received. Since agents can misreport their type, our aim is to design a mechanism that incentivises the agents to make truthful reports. We denote the reported type by $\hat{\theta}_i = \langle \hat{\mathbf{v}}_i, \hat{a}_i, \hat{d}_i, \hat{r}_i \rangle$. Our results use a common assumption in the online mechanism design literature (Hajiaghayi et al., 2005), that agents cannot report an earlier arrival or a later departure. In addition, we assume that agents cannot misreport a higher maximum consumption rate.

Assumption 2. Limited Misreports: *Agents cannot report an earlier arrival, a later departure, or a higher maximum consumption rate, i.e.,* $\hat{a}_i \geq a_i, \hat{d}_i \leq d_i, \hat{r}_i \leq r_i$ *must hold.*

In the following, *reports* that satisfy both Assumption 1 and Assumption 2 are said to be *admissible*. Given this assumption, our aim is to develop a mechanism which is dominant-strategy incentive compatible (DSIC), i.e., agents are best off reporting $\hat{\theta}_i = \theta_i$, no matter what other agents report.

Formally, let $\theta_I = \{\theta_i | i \in I(t)\}$ denote the types of all agents at time t, and $\theta_{-i} = \{\theta_j | j \in I(t), j \neq i\}$ the types of all agents except i, and similarly, $\hat{\theta}_I$ and $\hat{\theta}_{-i}$ denote the corresponding reported types. Note that, for brevity, we remove the dependence on t in this notation. Furthermore, $k_i^{\langle t \rangle}$ denotes the *endowment* (or number of units allocated so far) at the beginning of time t, not including the allocation at time t, and $\mathbf{k}^{\langle t \rangle} = \langle k_1^{\langle t \rangle}, k_2^{\langle t \rangle}, \ldots \rangle$. Furthermore, $k_i = k_i^{\langle \hat{d}+1 \rangle}$ denotes agent i's endowment upon its reported departure.

A mechanism is defined by an *allocation policy*, $\pi_i^{\langle t \rangle}(\hat{\theta}_I | \mathbf{k}^{\langle t \rangle}), i \in I(t)$, which determines the number of units allocated to agent i at time t given the current endowment, and a *payment policy*, $x_i(\hat{\theta}_{-i} | k_i), i \in I(t)$, which calculates the total payment for the allocated units. The allocation is made online, but the payment only needs be finalised upon the reported departure of an agent. Note that $k_i^{\langle t+1 \rangle} = k_i^{\langle t \rangle} + \pi_i^{\langle t \rangle}(\hat{\theta}_I | \mathbf{k}^{\langle t \rangle})$. We also use $\pi_i(\hat{\theta}_i) = \sum_{t=\hat{a}_i}^{\hat{d}_i} \pi_i^{\langle t \rangle}(\hat{\theta}_I | \mathbf{k}^{\langle t \rangle})$ to denote the total number of units allocated to agent i, given its reported type.

The aim is to find a mechanism satisfying the following property:

Definition 1. *(Dominant-Strategy Incentive Compatible (DSIC)) A mechanism is DSIC if reporting truthfully, i.e., $\hat{\theta}_i = \theta_i$, is a weakly dominant strategy. Formally, for all agents i, all admissible $\theta_i, \hat{\theta}_i \in \Theta$ and for all $\hat{\theta}_{-i}$:*

$$V(\pi_i(\theta_i), \theta_i) - x_i(\hat{\theta}_{-i}|\pi_i(\theta_i)) \geq V(\pi_i(\hat{\theta}_i), \theta_i) - x_i(\hat{\theta}_{-i}|\pi_i(\hat{\theta}_i)) \quad (1)$$

3.1 Application to Plug-In Hybrid Electric Vehicle Charging

In applying the model to the PHEV charging domain, agents compete for a limited charging capacity on behalf of their EV owners within a neighbourhood. We assume that the market for electricity for PHEV charging is separate from that for regular household consumption. Given this, the available supply, $S(t)$, for charging the PHEV is the residual supply once regular household consumption has been removed. The supply can also include electricity from uncertain sources, such as a shared renewable generator, e.g., a shared neighbourhood wind turbine or a solar panel installation.

In this scenario, a *unit* of electricity is defined as the amount of kWh when charging at the *lowest rate* during that interval (e.g., if the lowest rate is 6.5 A, then for 230 V and hourly slots, a unit is 6.5 A $\cdot$ 230 V $\cdot$ 1 h = 1.495 kWh). Charging points typically allow for the charging to occur at different rates, and so the maximum consumption rate, r_i, refers to the maximum charging rate of the charging point given the battery. Since units are indivisible, this means that the charging rate needs to be a multiple of the lowest rate. For example, if an agent is allocated 2 units in a single time step, then the charging rate is twice the lowest charging rate for that time interval. Given the focus on network capacity, supply is assumed perishable and capacity left unused at time t cannot be allocated later.

The time between arrival, a_i, and departure, d_i, refers to the interval where the vehicle is available for charging (i.e., it is at home and not being used). However, if the agent believes that it can benefit from delaying its arrival, then it can wait before plugging the vehicle into the electricity network. Therefore, the arrival report $\hat{a}_i$ is the time at which the owner physically plugs a vehicle into the electricity network, and a misreport consists of not plugging in on arrival. Similarly, the reported departure, $\hat{d}_i$, simply represents the time when the vehicle is unplugged from the electricity network. Although arrival and departure are modeled as part of the reported type, in practice these do not need to be communicated in advance to the mechanism, and are simply observed as they occur, i.e., when the owner plugs in or unplugs the vehicle.

The limited misreports assumption (Assumption 2) is reasonable in this context, since agents cannot physically plug in the EV if it is not at home. The requirement that $\hat{r}_i \leq r_i$ is also natural for PHEV charging. While most electric batteries can be configured to charge at a slower rate, charging them at a faster rate than the one allowed by the manufacturer might destroy them.

At the time of arrival, an agent needs to report its marginal valuation vector $\hat{\mathbf{v}}_i$. There is a clear interpretation of the marginal valuation because PHEVs can always use petrol as a substitute for electricity. The marginal value for an additional unit of charge is the expected money saved by not having to incur the cost of petrol.[2] In determining the exact value, we also need to consider the amount of purchased electricity the owner will use in

2. Note that we implicitly assume that petrol (i.e. gasoline) is always available as a substitute, and that there are sufficient refueling points so that the vehicle does not run out of petrol.

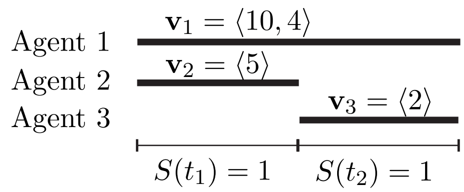

Figure 1: Example showing arrivals, departures, and valuation vectors of 3 agents.

expectation. For example, if she is certain to use the first unit of electricity, her maximum willingness to pay would be equal to the equivalent cost of petrol. As units become less likely to be used, their expected value decreases. The value of a marginal unit is the expected savings compared to using the petrol alternative. In Section 7 we provide a detailed analysis, confirming the non-increasing marginal value (Assumption 1) for hybrid EVs.

4. The Online Mechanism

In this section, we first present a simple greedy allocation policy, that we show cannot be coupled with a payment rule to provide truthfulness. Continuing, we combine it with two variations on the idea of modifying the allocation generated by the greedy rule, and provide a theoretical analysis of their properties.

4.1 Greedy Allocation Policy

Let the vector $\hat{\mathbf{v}}_i^{\langle t \rangle} = \langle \hat{v}_{i,k_i^{\langle t \rangle}+1}, \ldots, \hat{v}_{i,k_i^{\langle t \rangle}+\hat{r}_i} \rangle$ denote agent i's reported marginal values for the next $\hat{r}_i$ units, given its endowment $k_i^{\langle t \rangle}$ at time t. This is the agent's reported *willingness to pay* for any units available at time t and, in what follows, we refer to vector $\hat{\mathbf{v}}_i^{\langle t \rangle}$ as the *active* (reported) marginal valuations at time t. Furthermore, let $\mathcal{V}^{\langle t \rangle}$ denote the multiset of such values from all agents that are present in the market at time t, i.e., from all $i \in I(t)$ such that $\hat{a}_i \leq t \leq \hat{d}_i$. Next, we define a set operator $\max_{\langle k \rangle} \mathcal{V}$ to return the highest k elements of multiset $\mathcal{V}$ (or, if $|\mathcal{V}| < k$, to return $\mathcal{V}$). Then, the greedy allocation policy is:

Definition 2 (Greedy Allocation Policy)**.** *At each time step t, allocate the $S(t)$ units so that every agent receives one unit of the good for each of its active marginal valuations included in* $\max_{\langle S(t) \rangle} \mathcal{V}^{\langle t \rangle}$.

We ignore issues with tie breaking throughout this paper to simplify the exposition. We note, however, that all results presented hold when implementing either a random tie breaking rule, or a "first come, first served" rule, that breaks ties in favour of the agent that arrived in the market first.

To provide an example for the greedy allocation, consider the active marginal valuations given in Table 2. In this case, the multiset of active marginal valuations consists of $\mathcal{V}^{\langle t \rangle} = \langle 7, 6 \rangle \cup \langle 10, 6, 6 \rangle \cup \langle 8 \rangle = \langle 7, 6, 10, 6, 6, 8 \rangle$ (in no particular order). Then, if $S(t) = 3$, the highest active marginal values it will allocate the resources to are the agents with marginal values $10, 8$, and 7. Thus, in this example, each agent will receive 1 unit.

In order to show why greedy is not always DSIC, consider the example illustrated in Figure 1, involving 2 time steps and 3 agents. Suppose that supply is $S(t) = 1$ at each time

Supply, Agents and Preferences (Section 3)

$S(t)$	Supply of perishable units at time t
$I(t) = \{1, 2, ...\}$	The set of agents who have arrived so far at time t
$\theta_i = \langle \mathbf{v}_i, a_i, d_i, r_i \rangle$	Type of agent i
$\mathbf{v}_i = \langle v_{i,1}, v_{i,2}, ... \rangle$	Marginal valuation vector, where $v_{i,k}$ is the value for the k^{th} unit
a_i, d_i, r_i	Arrival time, departure time, and maximum consumption rate
Θ	Set of all admissible types, i.e., subject to Assumptions 1 and 2
$\theta_I = \{\theta_i \mid i \in I(t)\}$	The types of all agents $i \in I(t)$
$\theta_{-i} = \{\theta_j \mid j \in I(t), j \neq i\}$	The types of all agents $j \in I(t)$ except i
$\hat{\theta}, \hat{\theta}_I, \hat{\theta}_{-i}$	Reported types
$V(k, \theta_i) = \sum_{j=1}^{k} v_{i,j}$	Total value given k units are allocated to i between a_i and d_i
$U(k, x, \theta_i) = V(k, \theta_i) - x$	Agent utility, where x is the payment to the mechanism

General Mechanisms (Section 3)

$k_i^{\langle t \rangle}$	Agent i's endowment at the beginning of time t
$\mathbf{k}^{\langle t \rangle} = \langle k_1^{\langle t \rangle}, k_2^{\langle t \rangle}, ... \rangle$	The endowment of all agents $i \in I(t)$ at the beginning of time t
$k_i = k_i^{\langle \hat{d}+1 \rangle}$	Agent i's endowment on reported departure
$\pi_i^{\langle t \rangle}(\hat{\theta}_I \mid \mathbf{k}^{\langle t \rangle})$	Allocation policy, i.e., number of units allocated to i at time t given all agents' reports $\hat{\theta}_I$ and current endowments $\mathbf{k}^{\langle t \rangle}$
$x_i(\hat{\theta}_{-i} \mid k_i)$	Payment policy, i.e., agent i's payment given the reported types of other agents, and agent i's endowment on departure.

Greedy Policy and DSIC Mechanism (Sections 4 and 4.2)

$\hat{\mathbf{v}}_i^{\langle t \rangle} = \langle \hat{v}_{i,k_i^{\langle t \rangle}+1}, ..., \hat{v}_{i,k_i^{\langle t \rangle}+\hat{r}_i} \rangle$	Agent i's reported active marginal values
$\mathcal{V}^{\langle t \rangle} = \bigcup_{i \in I(t)} \hat{\mathbf{v}}_i^{\langle t \rangle}$	Multiset of reported active marginal values (where the union operator is used as a multiset operator throughout the paper)
$\mathcal{V}_{-i}^{\langle t \rangle}$	Active reported marginal values if agent i would never have been present in the market
$\mathcal{V}_{-i\cup 0}^{\langle t \rangle}$	Zeros are added to this multiset to ensure that $\lvert\mathcal{V}_{-i\cup 0}^{\langle t \rangle}\rvert \geq S(t)$
$\max\limits_{\langle k \rangle} \mathcal{V}, \min\limits_{\langle k \rangle} \mathcal{V}$	Returns the highest respectively lowest k elements of multiset $\mathcal{V}$
$\mathcal{E}_i^{\langle t \rangle} = \min\limits_{\langle r_i \rangle} \left(\max\limits_{\langle S(t) \rangle} \mathcal{V}_{-i\cup 0}^{\langle t \rangle} \right)$	Externality imposed by agent i on others, i.e., the marginal valuations of those that are missing out due to agent i being allocated $\min(r_i, S(t))$ units at time t
$\mathbf{p}_i^{\langle t \rangle} = \mathrm{incr} \left(\bigcup_{t'=\hat{a}_i}^{t} \mathcal{E}_i^{\langle t' \rangle} \right)$	Vector of marginal payments at time t, where $p_{i,k}^{\langle t \rangle}$ is the price that agent i is charged for the k^{th} unit, and the operator incr orders elements from a multiset in increasing order
$\mathbf{p}_i = \mathbf{p}_i^{\langle \hat{d}_i \rangle}$	Marginal payment vector on reported departure of agent i

Table 1: Main notation with references to sections where it is first introduced.

agent (i)	$\hat{r}_i$	$\hat{\mathbf{v}}_i$	$k_i^{\langle t \rangle}$	$\hat{\mathbf{v}}_i^{\langle t \rangle}$
1	2	$\langle 8,7,6,5 \rangle$	1	$\langle 7,6 \rangle$
2	3	$\langle 10,6,6,4,4 \rangle$	0	$\langle 10,6,6 \rangle$
3	1	$\langle 9,8,8,7 \rangle$	2	$\langle 8 \rangle$

Table 2: Example of three agents' active marginal valuations $\hat{\mathbf{v}}_i^{\langle t \rangle}$, given their marginal valuations $\hat{\mathbf{v}}_i$, endowments $k_i^{\langle t \rangle}$, and maximum consumption rates $\hat{r}_i$.

step and $r_i = 1$ for all agents i. Assuming truthful agents, greedy would then allocate both units to agent 1, because agent 1 has the highest active marginal value at both time steps ($\hat{\mathbf{v}}_1^{\langle 1 \rangle} = \langle 10 \rangle > \hat{\mathbf{v}}_2^{\langle 1 \rangle} = \langle 5 \rangle$, and $\hat{\mathbf{v}}_1^{\langle 2 \rangle} = \langle 4 \rangle > \hat{\mathbf{v}}_3^{\langle 2 \rangle} = \langle 2 \rangle$).

Now, consider the question of finding a payment policy that makes the greedy allocation policy DSIC. How much should agent 1 pay? To answer this, note that the payment for the unit allocated at time $t = 1$ has to be at least 5. Otherwise, if agent 1 were present in the market only at time $t = 1$ and had a valuation $v_{1,1} \in (5 - \epsilon, 5)$, it would have an incentive to misreport $\hat{v}_{1,1} > 5$ and still win. Similarly, the payment for the unit allocated at time $t = 2$ has to be at least 2. Thus, the *minimum* payment of agent 1 if allocated two units is $x_1(\hat{\theta}|\pi_1 = 2) = 7$.

On the other hand, how much should agent 1 pay if it were allocated only one unit instead? We argue no more than 2. Suppose, for contradiction, that the payment were set at some larger value $x_1(\hat{\theta}|\pi_1 = 1) = 2 + \epsilon$ (where $\epsilon > 0$). Then if the agent's first marginal value $v_{1,1}$ were instead $2 < v_{1,1} < 2 + \epsilon$ (with remaining marginal values zero), then it would win in period 2, but it would pay $2 + \epsilon$ and hence have negative utility. However, if $x_1(\hat{\theta}|\pi_1 = 2) \geq 7$ and $x_1(\hat{\theta}|\pi_1 = 1) \leq 2$, then agent 1 wants only one unit, not two, as allocated by the greedy mechanism (its utility for one unit is greater than for two, as $10 - 2 > 10 + 4 - 7$). Hence, the greedy allocation policy cannot be made DSIC by setting payments.

The above example shows that the problem is not with a particular payment policy, but is intrinsic to the greedy allocation policy. In particular, the problem is that the allocation policy does not satisfy the necessary property of W-MON (Bikhchandani, Chatterji, Lavi, Mu'alem, Nisan, & Sen, 2006):

Definition 3 (Weak Monotonicity (W-MON))**.** *An allocation policy π, is W-MON if, for every $i \in I(t)$, $\theta_i = \langle \mathbf{v}_i, a_i, d_i, r_i \rangle, \theta'_i = \langle \mathbf{v}'_i, a'_i, d'_i, r'_i \rangle \in \Theta$, $\theta_{-i} \in \Theta^{N-1}$, where Θ is the set of all types* subject to non-increasing marginal valuations*, the following equation holds:*

$$V(\pi_i(\theta'_i), \theta'_i) - V(\pi_i(\theta_i), \theta'_i) \geq V(\pi_i(\theta'_i), \theta_i) - V(\pi_i(\theta_i), \theta_i) \quad (2)$$

In words, if changing an agent i's type (while keeping the types of other agents fixed) from a type θ_i to another type θ'_i changes the allocation of i from $\pi_i(\theta_i)$ to $\pi_i(\theta'_i)$, then the resulting difference in utilities of the new and original outcomes evaluated at the new type of agent i (denoted by function $V(\cdot, \theta'_i)$) must be no less than this difference in utilities evaluated at the original type of agent i (denoted by function $V(\cdot, \theta_i)$). Using this notion we will now demonstrate that the greedy allocation policy is not DSIC in our setting.

Theorem 1. *The greedy allocation policy is not DSIC in multi-valued domains with non-increasing marginal valuations.*

Proof. From Bikhchandani et al. (2006, Lemma 1) we know that a *necessary* condition for any DSIC allocation policy is that it should satisfy W-MON (see Definition 3). If Equation 2 is true, then it can be shown that the following must hold:

$$\text{if } \pi_i(\theta_i') > \pi_i(\theta_i) \text{ then } \sum_{k=\pi_i(\theta_i)+1}^{\pi_i(\theta_i')} v_{i,k}' \geq \sum_{k=\pi_i(\theta_i)+1}^{\pi_i(\theta_i')} v_{i,k} \tag{3}$$

In words, more units should only be allocated to one type compared to another type, if that type also has higher marginal values for those units. Consider again the example in Figure 1, where we look at the W-MON condition for agent 1 by varying its type. We keep the arrival and departure fixed, i.e., $a_i' = a_i, d_i' = d_i$. Suppose that $v_1' = \langle 10, 4 \rangle$ as in the example, but $v_1 = \langle 4+\epsilon, 4+\epsilon \rangle$, where $0 < \epsilon < 1$. Note that $\pi_1(\theta_1') = 2$, and $\pi(\theta_1) = 1$ (agent 1 is not allocated the unit in the first time step if its type were changed from θ_1' to θ_1). Since θ_1' is allocated an additional unit (compared to θ_1), W-MON requires that θ_1' values the second unit higher or equal to θ_1. However, we can see that $v_{1,2}' < v_{1,2}$, thereby violating Equation 3. This example demonstrates that the greedy allocation policy is not W-MON, and therefore not DSIC. □

4.2 Achieving Truthfulness Through Cancellation

Addressing the problem with W-MON, we consider two types of modifications to the allocation decision of the greedy policy, both designed to achieve monotonicity. The first is *immediate* cancellation, where units are simply left unallocated, i.e., none of the agents receive the unit, even if there is a demand. The second is *on-departure* cancellation, where units are initially allocated using the greedy approach, but then on departure of the agent, any overallocated units are removed.[3]

The model with on-departure cancellation is more efficient because it generally requires fewer cancellations, and is also more computationally efficient in calculating the payments and allocations. However, depending on the domain, it may not always be possible to remove units once they are allocated.

In the following, we detail the allocation policies and explain how payments are computed. We then give an example that shows the difference between the two mechanisms and, lastly, we provide an analysis of the economic properties of the mechanisms.

In defining the allocation policies, we show how we can compute an agent's *marginal payment vector*, which determines, for each additional unit, the price the agent would need to pay for that unit. These marginal payments are then used to determine both when to cancel an allocation, as well as an agent's total payment for a given allocation.

3. In the PHEV setting, this corresponds to first charging the battery and later discharging any overallocated units.

A necessary condition for truthfulness is that payments cannot depend on an agent's report except on the effect it has on the allocation; e.g., see the work of Nisan, Roughgarden, Tardos, and Vazirani (2007, Proposition 9.27). To this end, let $\mathcal{V}_{-i}^{\langle t \rangle}$ denote the multiset of the active marginal valuations of all agents in the market at time t, if agent i were removed and the market were rerun from $\hat{a}_i \leq t$ onwards.

We cannot simply derive $\mathcal{V}_{-i}^{\langle t \rangle}$ from $\mathcal{V}^{\langle t \rangle}$ since removing agent i could affect the endowments of the other agents at earlier time steps. For example, for the setting from Table 2, both agents 1 and 3 have endowments at time t and therefore removing either of these agents is likely to increase the endowments of the other agents, thereby changing the dynamics of the entire market. To ensure that $\mathcal{V}_{-i}^{\langle t \rangle}$ is truly independent of agent i, the market needs to be re-run from the very point when agent i first entered the market, and this process needs to be repeated for each agent $i \in I(t)$ such that $\hat{a}_i \leq t \leq \hat{d}_i$, whose payments we are computing.

In case $|\mathcal{V}_{-i}^{\langle t \rangle}| < S(t)$, we furthermore add a number of zero-valued bids and refer to this enlarged set as $\mathcal{V}_{-i\cup 0}^{\langle t \rangle}$, to ensure that $|\mathcal{V}_{-i\cup 0}^{\langle t \rangle}| \geq S(t)$. Next, similar to the operator $\max_{\langle k \rangle}$, we define the set operator $\min_{\langle k \rangle} \mathcal{V}$ to return the *lowest* k elements of multiset $\mathcal{V}$ (or, if $|\mathcal{V}| < k$, to return $\mathcal{V}$). We define the *externality* that agent i would impose on other agents if it won $\min(r_i, S(t))$ out of $S(t)$ units at time t as:

$$\mathcal{E}_i^{\langle t \rangle} = \min_{\langle r_i \rangle} \left(\max_{\langle S(t) \rangle} \mathcal{V}_{-i\cup 0}^{\langle t \rangle} \right)$$

Note that the cardinality of $\mathcal{E}_i^{\langle t \rangle}$ is equal to $|\mathcal{E}_i^{\langle t \rangle}| = \min(r_i, S(t))$. Intuitively, the multiset $\mathcal{E}_i^{\langle t \rangle}$ contains the marginal values from other agents that would lose out if agent i were to win r_i units at time t. For example, let $\mathcal{V}_{-1}^{\langle t \rangle} = \langle 1, 4, 5, 7, 9, 10 \rangle$ (sorted for convenience), $S(t) = 4$, and $r_1 = 2$. Then, $\mathcal{E}_1^{\langle t \rangle} = \langle 5, 7 \rangle$.

If an agent were active for only a single time step, then the externality would specify the payment for each unit. That is, using the same example, if agent 1 were allocated a single unit by the mechanism, its payment would be 5. On the other hand, if it were allocated 2 units, its payment would be $5 + 7 = 12$. The intuition here is the same as in the regular Vickrey-Clarke-Groves (VCG) mechanism, because the total payment corresponds to the sum of the externalities.

To compute the overall payments online, we need to combine these externalities across all time steps in the agent's active period up to current time t. To do this, we define an ordered vector of *marginal payments*, $\mathbf{p}_i^{\langle t \rangle}$, as follows:

$$\mathbf{p}_i^{\langle t \rangle} = \text{incr}\left(\bigcup_{t'=\hat{a}_i}^{t} \mathcal{E}_i^{\langle t' \rangle} \right),$$

where 'incr' is an operator that orders elements from a multiset in increasing order, and we use the union symbol to denote the union of multisets (and so the same element can appear multiple times).

Now, $p_{i,k}^{\langle t \rangle}$ is the price that agent i is charged for the k^{th} unit of the good. Intuitively, this is the minimum valuation that the agent could report for winning this unit by time t. These prices are adjusted in each time step. In particular, since the vector is in increasing order, and elements are only added as time increases, $p_{i,k}^{\langle t\rangle}$ either stays the same or *decreases* for a given k, but can never increase. In the following, we use $p_{i,k}$ to denote agent i's marginal payment of the k^{th} unit at time $\hat{d}_i$.

Given this, the decision and payment policies are defined as follows:

- **Allocation Policy** The decision to allocate consists of two stages:

  Stage 1 At each time step t, *pre-allocate* using the greedy allocation policy (see Definition 2).

  Stage 2 We consider two variations of how to decide to cancel the pre-allocation:

  - **Immediate Cancellation (`IM`).** Leave any unit unallocated whenever the marginal payment *at time* t for this unit is greater than the marginal value, i.e., whenever:

    $$\hat{v}_{i,k} < p_{i,k}^{\langle t\rangle} \text{ for } k_i^{\langle t\rangle} < k \leq k_i^{\langle t\rangle} + \pi_i^{\langle t\rangle}$$

  - **On-Departure Cancellation (`OD`).** For each departing agent, cancel the allocation of any unit $k \leq k_i$ where $\hat{v}_{i,k} < p_{i,k}$.

- **Payment Policy** Payment always occurs on reported departure. Given that k_i units are allocated to agent i, the payment collected from i is:

$$x_i(\hat{\theta}_{-i}|k_i) = \sum\nolimits_{k=1}^{k_i} p_{i,k} \tag{4}$$

In the following, we refer to our two mechanisms with immediate and on-departure cancellation by `IM` and `OD` respectively, and the corresponding allocation policies by π^{im} and π^{od}. If no distinction between the two mechanisms is made, π is used.

The payment policy mirrors the allocation policy. For example, if immediate cancellation is used, then for each agent i and for all times t, the values of the $\mathbf{p}_i^{\langle t\rangle}$ vector are computed by re-running the market, in the absence of agent i using immediate cancellation, based on the reports of the other agents. Conversely, if on-departure cancellation is used, the same policy should be used in computing the $\mathbf{p}_i$ prices.

4.3 Examples

To demonstrate how the two mechanisms work, we present two examples. The aim of the first example is to show the difference between immediate and on-departure cancellation. The second example illustrates the effect of changing the maximum consumption rates. The example shows that in specific instances, increasing an agent's maximum consumption rate can actually increase the number of cancellations.

4.3.1 Example 1

The first example extends the setting shown in Figure 1 to include a third time step, $t = 3$. Both agents 1 and 3 remain in the market at $t = 3$ (i.e., $d_1 = d_3 = 3$) and no new agents arrive. Furthermore, $S(t) = 1$ in $t \in \{1, 2, 3\}$, and so there are now three units to be allocated in total. As before, we suppose that $r_i = 1$ for all agents. Table 3 shows the

	agent 1: $a_1 = 1$, $d_1 = 3$, $\mathbf{v}_1 = \langle 10, 4\rangle$, $r_1 = 1$	**agent 2:** $a_2 = 1$, $d_2 = 1$, $\mathbf{v}_2 = \langle 5\rangle$, $r_2 = 1$	**agent 3:** $a_3 = 2$, $d_3 = 3$, $\mathbf{v}_3 = \langle 2\rangle$, $r_3 = 1$
$t = 1$	$k_1^{\langle 1\rangle} = 0$, $\mathbf{v}_1^{\langle 1\rangle} = \langle 10\rangle$ $\mathcal{V}_{-1}^{\langle 1\rangle} = \mathcal{E}_1^{\langle 1\rangle} = \langle 5\rangle$ $\mathbf{p}_1^{\langle 1\rangle} = \langle 5\rangle$ $\pi_1^{im\langle 1\rangle} = \pi_1^{od\langle 1\rangle} = 1$	$k_2^{\langle 1\rangle} = 0$, $\mathbf{v}_2^{\langle 1\rangle} = \langle 5\rangle$ $\mathcal{V}_{-2}^{\langle 1\rangle} = \mathcal{E}_2^{\langle 1\rangle} = \langle 10\rangle$ $\mathbf{p}_2^{\langle 1\rangle} = \langle 10\rangle$ $\pi_2^{im\langle 1\rangle} = \pi_2^{od\langle 1\rangle} = 0$	
$t = 2$	$k_1^{\langle 2\rangle} = 1$, $\mathbf{v}_1^{\langle 1\rangle} = \langle 4\rangle$ $\mathcal{V}_{-1}^{\langle 2\rangle} = \mathcal{E}_1^{\langle 2\rangle} = \langle 2\rangle$ $\mathbf{p}_1^{\langle 2\rangle} = \langle 2, 5\rangle$ $\pi_1^{im\langle 2\rangle} = 0$ $\pi_1^{od\langle 2\rangle} = 1$		$k_3^{\langle 2\rangle} = 0$, $\mathbf{v}_3^{\langle 1\rangle} = \langle 2\rangle$ $\mathcal{V}_{-3}^{\langle 2\rangle} = \mathcal{E}_3^{\langle 2\rangle} = \langle 4\rangle$ $\mathbf{p}_3^{\langle 2\rangle} = \langle 4\rangle$ $\pi_3^{im\langle 2\rangle} = \pi_3^{od\langle 2\rangle} = 0$
$t = 3$ `IM`	$k_1^{\langle 3\rangle} = 1$, $\mathbf{v}_1^{\langle 1\rangle} = \langle 4\rangle$ $\mathcal{V}_{-1\cup 0}^{\langle 3\rangle} = \mathcal{E}_1^{\langle 3\rangle} = \langle 0\rangle$ $\mathbf{p}_1^{\langle 3\rangle} = \langle 0, 2, 5\rangle$ $\pi_1^{im\langle 3\rangle} = 1$		$k_3^{\langle 3\rangle} = 0$, $\mathbf{v}_3^{\langle 1\rangle} = \langle 2\rangle$ $\mathcal{V}_{-3}^{\langle 3\rangle} = \mathcal{E}_3^{\langle 3\rangle} = \langle 4\rangle$ $\mathbf{p}_3^{\langle 3\rangle} = \langle 4, 4\rangle$ $\pi_3^{im\langle 3\rangle} = 0$
$t = 3$ `OD`	$k_1^{\langle 3\rangle} = 2$, $\mathbf{v}_3^{\langle 1\rangle} = \langle\rangle$ $\mathcal{V}_{-1\cup 0}^{\langle 3\rangle} = \mathcal{E}_1^{\langle 3\rangle} = \langle 0\rangle$ $\mathbf{p}_1^{\langle 3\rangle} = \langle 0, 2, 5\rangle$ $\pi_1^{od\langle 3\rangle} = 0$		$k_3^{\langle 3\rangle} = 0$, $\mathbf{v}_3^{\langle 1\rangle} = \langle 2\rangle$ $\mathcal{V}_{-3\cup 0}^{\langle 3\rangle} = \mathcal{E}_3^{\langle 3\rangle} = \langle 0\rangle$ $\mathbf{p}_3^{\langle 3\rangle} = \langle 0, 4\rangle$ $\pi_3^{od\langle 3\rangle} = 1$

Table 3: Example run of the mechanism with 3 agents and 3 time steps for the `IM` and `OD` mechanisms. Grey cells indicate different values for `IM` and `OD` policies.

endowments $k_i^{\langle t\rangle}$, the active marginal valuations $\mathcal{V}_{-i}^{\langle t\rangle}$, the externalities, $\mathcal{E}_i^{\langle t\rangle}$, the marginal payments $\mathbf{p}_i^{\langle t\rangle}$, and the allocation decisions $\pi_i^{\langle t\rangle}$ at different time steps.

We start by considering the allocations and payments using *immediate cancellation*. At time $t = 1$, Stage 1 of the mechanism pre-allocates the unit to agent 1, and since $v_{1,1} = 10 \geq p_{1,1}^{\langle 1\rangle} = 5$, this pre-allocation is not cancelled in the second stage. At time $t = 2$, the unit again gets pre-allocated to agent 1 since its active marginal value is greater than that of agent 3, i.e., $\mathbf{v}_1^{\langle 2\rangle} = \langle 4\rangle > \mathbf{v}_3^{\langle 2\rangle} = \langle 2\rangle$. However, $\mathcal{V}_{-1}^{\langle 2\rangle} = \langle 2\rangle$ is inserted at the beginning of the $\mathbf{p}_1^{\langle 2\rangle}$ vector, and as a result $v_{1,2} = 4 < p_{1,2}^{\langle 2\rangle} = 5$ (at these prices, agent 1 prefers to be allocated one unit instead of two). Consequently, this pre-allocation gets cancelled and the unit goes to neither of the agents.

At time $t = 3$, the active marginal value of agent 1 is still $\mathbf{v}_1^{\langle 3\rangle} = \langle 4\rangle$, since its endowment is unchanged, and since agent 1 still has the highest active marginal value, it is again pre-allocated the unit. To calculate the marginal payment of agent 1, recall that the allocation policy needs to be recomputed with agent 1 entirely removed from the market. In that case agent 3 would have been allocated a unit at time $t = 2$, and thus at time $t = 3$ the active marginal value of this agent is 0. Thus, the value of 0 is inserted in the $\mathbf{p}_1^{\langle 3\rangle}$ vector. At $t = 3$, however, now $v_{1,2} = 4 \geq \mathbf{p}_{1,2}^{\langle 3\rangle} = 2$, and therefore the pre-allocation is not cancelled.

An interesting exercise is to see what happens to the marginal payment vector of agent 3 at $t = 3$. To calculate this, we have to remove agent 3 and rerun the market from $t = 2$. In

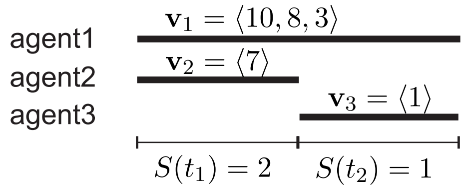

Figure 2: Example showing arrivals, departures, and valuation vectors of 3 agents.

this case, at time $t = 2$, the marginal payment vector of agent 1 becomes $\mathbf{p}_1^{\langle 2 \rangle} = \langle 0, 5 \rangle$. Since the marginal payment for the second unit, $p_{1,2}^{\langle 2 \rangle}$, remains unchanged, the pre-allocation is still cancelled! Therefore, even when agent 3 is not in the market, the second unit remains unallocated, and agent 1's active marginal value at $t = 3$ is again $\mathbf{v}_1^{\langle 3 \rangle} = \langle 4 \rangle$, and so $\mathcal{V}_{-3}^{\langle 3 \rangle} = \langle 4 \rangle$. Given this, agent 3's marginal payments become $\mathbf{p}_3^{\langle 3 \rangle} = \langle 4, 4 \rangle$. Note that the marginal payment for the first unit (4) is higher than the marginal value for this unit (2), and this is consistent with the allocation. Otherwise, if the marginal payment had been lower, agent 3 would have an incentive to overreport and 'win' the unit at $t = 3$.

So, in case of immediate cancellation, two out of three units are allocated to agent 1, and that agent pays $x_1 = p_{1,1}^{\langle 3 \rangle} + p_{1,2}^{\langle 3 \rangle} = 0 + 2 = 2$. The third unit is not allocated to any agent. Note that this unit cannot go to agent 3, because the payment would have been $p_{3,1}^{\langle 3 \rangle} = 4$, resulting in a negative utility for agent 3.

Now consider the same setting but with *on-departure cancellation*. The first two time steps are as before, except that there is no cancellation at $t = 2$ (since this will be done on departure if needed). This changes the endowment state of agent 1 at $t = 3$, and therefore the active marginal value of agent 1 at $t = 3$ is equal to $\mathbf{v}_1^{\langle 3 \rangle} = \langle \rangle$, and this is the same if agent 3 is removed from the market. Therefore, the unit is pre-allocated to agent 3, and the payment for this unit is $p_{3,1} = p_{3,1}^{\langle 3 \rangle} = 0$. The vector $\mathbf{p}_1^{\langle 3 \rangle}$ remains unchanged compared to the immediate case. At this point, there is no longer a need to cancel one of the pre-allocations of agent 1, since it has received $k = 2$ units, the same allocation as with the immediate cancellation policy, and note that $v_{1,2} > p_{1,2}$.

No pre-allocations are cancelled with the on-departure policy, and so this policy is more efficient. As we will show in the remainder of the paper, the on-departure policy never cancels more, and typically cancels fewer pre-allocations compared to the immediate one. Still, it is possible to construct examples where, in the worst case, half of the pre-allocations need to be cancelled, even with the on-departure policy.

Furthermore, note that some units are given away for free (i.e., the payment for these units is zero). This is a standard problem with auctions if there is insufficient competition, and can be trivially resolved by e.g. introducing a minimum price or reserve price for each unit of the good. However, this will reduce efficiency since units will remain unallocated if they fall below the reserve price. We do not consider reserve prices in this paper, but the economic properties of the mechanism continue to hold with reserve prices, as long as the reserve prices are the same for all time points (otherwise there could be an incentive to misreport the arrival time).

4.3.2 Example 2

The next example, depicted in Figure 2, again shows a setting with two time steps and three agents, but with different preferences and now the supply in the first step is two units, and we change the maximum consumption rate of agent 1. We consider two cases:

The maximum consumption rate of agent 1 is $r_1 = 1$.[4] In this case, at most one marginal value is taken from each agent per time step. At time t_1, marginal valuations $v_{1,1} = 10$ of agent 1, and $v_{2,1} = 7$ of agent 2 are pre-allocated, while at time t_2, marginal value $v_{1,2} = 8$ of agent 1 is pre-allocated. The prices charged to agent 1 are: $p_1 = \langle 0, 1\rangle$, because without agent 1 in the market, there would be a free, spare unit at time t_1 and the available unit at t_2 would sell to agent 3 for 1. No pre-allocation gets cancelled in this case, and the actual allocation is equivalent to the optimal offline allocation.

The maximum consumption rate of agent 1 is $r_1 = 2$. Then, at time t_1, the greedy policy described above allocates the two marginal values of agent 1: $v_{1,2} = 10$ and $v_{1,2} = 8$, as they are both higher than $v_{2,1} = 7$, and agent 2 drops out of the market. At time t_2, the unit is again pre-allocated to agent 1 (due to the marginal value of 3 being higher than 1). However, now the marginal payments vector required from agent 1 is $p_1 = \langle 0, 1, 7\rangle$, while the marginal valuations are $v_1 = \langle 10, 8, 3\rangle$. Given the prices, agent 1 prefers two units to three (because $10 + 8 - 1 > 10 + 8 + 3 - 1 - 7$), so the third is cancelled. The overall efficiency is lower, as the pre-allocation of the third available unit is now cancelled, whereas with $r_1 = 1$ it was allocated to agent 2. Note, however, that although the efficiency is much lower, agent 1 has an incentive to declare its true maximum consumption rate $r_1 = 2$ as, in this case, its payment does not change.

4.4 Establishing Truthfulness

In this section we prove that the above mechanisms are DSIC under the assumptions of non-increasing marginal valuations (Assumption 1) and limited misreports (Assumption 2).

In the following, we will first establish DSIC with respect to valuations only, and prove truthful reporting of arrival and departure times separately. In more detail, we proceed in the following 3 stages:

(i) We define the concept of a threshold policy, and show that, when coupled with an appropriate payment policy, and given any admissible pair $\langle \hat{a}_i, \hat{d}_i\rangle$, if an allocation policy is a threshold policy, then the mechanism is DSIC with respect to the valuations (Lemma 1).

(ii) We show that our allocation policy is a threshold policy (Lemma 2).

(iii) We show that if agents truthfully report their valuations, reporting $\hat{a}_i = a_i$, $\hat{d}_i = d_i$, and $\hat{r}_i = r_i$ is a weakly dominant strategy (Lemma 3).

These results are combined in Theorem 2 to show that our mechanism is DSIC.

Definition 4 (Threshold Policy)**.** *An allocation policy π is a threshold policy if, for a given agent i with fixed $\langle \hat{a}_i, \hat{d}_i, \hat{r}_i\rangle$ and $\hat{\theta}_{-i}$, there exists a marginally non-decreasing threshold vector τ, independent from the report $\hat{\mathbf{v}}_i$ made by agent i, such that following holds: $\forall k, \hat{\mathbf{v}}_i$: $\pi_i(\hat{\theta}_i, \hat{\theta}_{-i}) \geq k$ if and only if $\hat{v}_{i,k} \geq \tau_k$.*

4. Note that the other two agents only desire one unit, so their maximum consumption rate is irrelevant in this example.

In other words, a threshold policy has a potentially different threshold τ_k for each k, such that agent i will receive at least k units if and only if its reported valuation for the k^{th} item is at least τ_k.

A threshold policy satisfies W-MON, which is sufficient for DSIC in this domain since we have bounded agent valuations and the domain is completely ordered, meaning that all payoff types agree on the same weak preference ordering on all allocations (i.e., more is always weakly better than less), and indifference to the way goods are allocated to other agents (Bikhchandani et al., 2006). We show that our allocation policy has the threshold property, and thus satisfies W-MON, and that it also handles misreports of arrivals, departures and maximum charging rates.

Importantly, the vector τ has to be non-decreasing, i.e., $\tau_{k+1} \geq \tau_k$, and should be independent of the reported valuation vector $\hat{\mathbf{v}}_i$. Both of these properties are satisfied by the $\mathbf{p}_i$ vector, and we will use this to show that our mechanism is a threshold policy.

First, however, we show that a threshold policy with appropriate payments is DSIC with respect to the valuations:

Lemma 1. *Fixing admissible $\langle \hat{a}_j, \hat{d}_j, \hat{r}_j \rangle$ for all $j \in I$ and fixing $\hat{\theta}_{-i}$, if π is a threshold policy coupled with a payment policy:*

$$x_i(\hat{\theta}_i, \hat{\theta}_{-i}) = \textstyle\sum_{k=1}^{\pi_i(\hat{\theta}_i, \hat{\theta}_{-i})} \tau_k,$$

then if $\mathbf{v}_i$ is marginally non-increasing, reporting $\mathbf{v}_i$ truthfully is a weakly dominant strategy.

Proof. Agent i's utility can be rewritten as:

$$u_i(\hat{\theta}_i; \theta_i) = \textstyle\sum_{k=1}^{\pi_i(\hat{\theta}_i, \hat{\theta}_{-i})} (v_{i,k} - \tau_k)$$

Since τ is independent of $\hat{\mathbf{v}}_i$, agent i can only potentially benefit by changing the allocation, $\pi_i(\hat{\theta}_i, \hat{\theta}_{-i})$. Since the values of $\tau_{k+1} \geq \tau_k$ (non-decreasing threshold vector) and $v_{i,k+1} \leq v_{i,k}$ (non-increasing marginal values), by Definition 4 we have $v_{i,k} - \tau_k \geq 0$ for any $k \leq \pi_i(\theta_i)$ and $v_{i,k} - \tau_k \geq 0$ for any $k > \pi_i(\theta_i)$. Suppose that, by misreporting agent i is allocated $\pi_i(\hat{\theta}_i) > \pi_i(\theta_i)$, then $u_i(\hat{\theta}_i; \theta_i) < u_i(\theta_i; \theta_i)$ since:

$$\textstyle\sum_{k=\pi_i(\theta_i, \hat{\theta}_{-i})+1}^{\pi_i(\hat{\theta}_i, \hat{\theta}_{-i})} (v_{i,k} - \tau_k) < 0$$

Similarly, misreporting such that $\pi_i(\hat{\theta}_i, \hat{\theta}_{-i}) < \pi_i(\theta_i, \hat{\theta}_{-i})$ results in $u_i(\hat{\theta}_i; \theta_i) < u_i(\theta_i; \theta_i)$ since:

$$\textstyle\sum_{k=\pi_i(\hat{\theta}_i, \hat{\theta}_{-i})+1}^{\pi_i(\theta_i, \hat{\theta}_{-i})} (v_{i,k} - \tau_k) \geq 0$$

If misreporting has no effect on the allocation, the utility remains the same. Therefore, there is no incentive for agent i to misreport its valuations. □

Note that the greedy allocation policy is not a threshold policy. Indeed, we have shown already that it does not satisfy W-MON. The next lemma shows that the threshold condition holds if we cancel some allocations according to our policies, and if we set the threshold values to $\tau_k = p_{i,k}$.

Lemma 2. *Given non-increasing marginal valuations, the allocation policy π in Section 4.2 is (for either cancellation policy) a threshold policy where $\tau_k = p_{i,k}$.*

Proof. First, from the definition of vector $\mathbf{p}_i^{\langle t \rangle}$ and $\mathbf{p}_i$ from Section 4.2, the values of $\mathbf{p}_i^{\langle t \rangle}$ are independent of the reports $\hat{\mathbf{v}}_i$ made by agent i. This is because each of its component values $\mathcal{V}_{-i}^{\langle \hat{a}_i \rangle}, \dots, \mathcal{V}_{-i}^{\langle t \rangle}$ are computed based only on the reports of the other agents, by first removing agent i from the market. Note that $\mathbf{p}_i^{\langle t \rangle}$ and $\mathbf{p}_i$ are also affected by the reported arrival time, departure time, and maximum consumption rate of agent i, but in this lemma we are only concerned with truthful reporting of the agent's valuations, and take the reported arrival and deadline as given, and do not require these to be truthful at this point.

Second, we need to show two inequalities, thus the proof is done in two parts. **Part 1:** Whenever $v_{i,k} \geq p_{i,k}$, π_i allocates *at least* k units to agent i. **Part 2:** Whenever $v_{i,k} < p_{i,k}$, π_i allocates *strictly fewer* than k units to agent i.

Part 1: Let $v_{i,k} \geq p_{i,k}$. Suppose that agent i has uniform marginal values, $v_{i,k}$, for the first k units (i.e., $v_{i,1} = v_{i,2} = \ldots = v_{i,k}$). Note that externality $\mathcal{E}_i^{\langle t \rangle}$ contains the marginal values at time t that agent i will displace when winning up to r_i units of the good in time step t (that is, these marginal values will reappear in the next time step if the same agents remain in the market). Given this, as long as agent i has fewer than k units then, in Stage 1 of the mechanism, and at each time step that agent i is in the market, it will win exactly those units where the marginal values in $\mathcal{E}_i^{\langle t \rangle}$ are less than $v_{i,k}$, i.e. it will win all units $1 \leq j \leq |\mathcal{E}_i^{\langle t \rangle}|$ where $v_{i,k} \geq \mathcal{E}_{i,j}^{\langle t \rangle}$ (ignoring tie breaking). Note that the externalities (and thus the marginal payments) are calculated by removing agent i from the market from the very first time it entered, and so do not contain any displaced marginal values. However, even when, by winning a unit, agent i displaces the losing marginal value to a future time step, since this value is less or equal to $v_{i,k}$, it will not affect the allocations of the first k units in future time steps for agent i since it will continue to have a higher marginal value. Now, because $p_{i,j} \leq p_{i,k}$ for $j \leq k$ (by definition), there must be at least k units for which $p_{i,k} \geq \mathcal{E}_{i,j}^{\langle t \rangle}$, $1 \leq j \leq |\mathcal{E}_i^{\langle t \rangle}|$ in the period $\hat{a}_i \leq t \leq \hat{d}_i$, and since $v_{i,k} \geq p_{i,k}$, agent i wins at least k units in Stage 1.

Furthermore, whenever j units are won at a particular time step, the marginal payments for those units appear as within the first $k' + j$ first elements of the $\mathbf{p}_i^t$ vector, where k' is the number of units won at earlier time steps (since these are the values with the the lowest clearing payment, and they are ordered ascendingly). Because agent i wins a unit with externality $\mathcal{E}_{i,j}^{\langle t \rangle}$ in Stage 1 if and only if $v_{i,k} \geq \mathcal{E}_{i,j}^{\langle t \rangle}$ (given uniform valuations), it follows that $v_{i,k} = v_{i,j} \geq p_{i,j}$ whenever it wins a unit in Stage 1. Therefore, no pre-allocations are cancelled in Stage 2.

The above holds if agent i has uniform marginal values of $v_{i,k}$ for the first k units. In fact, however, because of non-increasing valuations, we have $v_{i,j} \geq v_{i,k}$, for all $1 \leq j \leq k$, and thus the allocation policy will allocate *at least* k units to agent i.

Part 2: Let $v_{i,k} < p_{i,k}$. First consider the *on-departure cancellation* case. As per the definition of Stage 2 of the mechanism, the allocation of unit k is cancelled. However, we still need to show that any pre-allocated units $j > k$ are cancelled as well. Since $p_{i,j} \geq p_{i,k}$ (by definition) and $v_{i,j} \leq v_{i,k}$ (since valuations are marginally non-increasing by assumption) for all $j > k$, it follows that $v_{i,j} < p_{i,j}$ for all $j > k$. Therefore even if Stage 1 pre-allocates

k or more units, these will be cancelled in Stage 2, and thus strictly fewer than k units remain.

Now consider the *immediate cancellation* case. Note that $p_{i,k} \leq p_{i,k}^{\langle t \rangle}$ for $t_k \leq t \leq d_i$, where t_k is the time when the k^{th} unit was allocated. That is, marginal payment values can only decrease over time. Since $v_{i,k} < p_{i,k}$ (by assumption) and $p_{i,k} \leq p_{i,k}^{\langle t_k \rangle}$, it follows that $v_{i,k} \leq p_{i,k}^{\langle t_k \rangle}$. Thus it follows that $v_{i,k} < p_{i,k}^{\langle t \rangle}$ for any $(a_i + k - 1) \leq t \leq d_i$. Consider a case where, at time t_k, the k^{th} unit is allocated in Stage 1. As a result, pre-allocation of the k^{th} unit will always be cancelled at time t_k in Stage 2 of the allocation policy. Therefore, the final result is an allocation of strictly fewer than k units. □

By setting $\tau_k = p_{i,k}$, the payment policy in Equation 4 corresponds to the payment policy in Lemma 1. Therefore the proposed mechanism is shown to be DSIC in valuations.

We complete the proof by showing that truthful reporting of the arrival and departure times are also DSIC given limited misreports, now assuming truthful reporting of $\mathbf{v}_i$.

Lemma 3. *Given limited misreports, and assuming that truthfully reporting $\hat{\mathbf{v}}_i = \mathbf{v}_i$ is a dominant strategy for any given arrival,departure and maximum consumption rate reports $\langle \hat{a}_i, \hat{d}_i, \hat{r}_i \rangle$, then (subject to limited misreports) it is a dominant strategy to report $\hat{a}_i = a_i$, $\hat{d}_i = d_i$, and $\hat{r}_i = r_i$.*

Proof. Let $\mathbf{p}_i^{\langle \hat{a}_i, \hat{d}_i, \hat{r}_i \rangle}$ denote the vector of increasingly ordered marginal clearing values (computed without i), given the agent reports $\hat{\theta}_i = \langle \mathbf{v}_i, \hat{a}_i, \hat{d}_i, \hat{r}_i \rangle$. By reporting type $\hat{\theta}_i$, the agent is allocated $\pi_i(\hat{\theta}_i)$ items, and its total payment is: $\sum_{j=1}^{\pi_i(\hat{\theta}_i)} p_{i,j}^{\langle \hat{a}_i, \hat{d}_i, \hat{r}_i \rangle}$. For each agent i, misreporting from θ_i to $\hat{\theta}_i$ results in one of two cases:

$\boldsymbol{\pi_i(\hat{\theta}_i) = \pi_i(\theta_i)}$: Misreporting by agent i cannot change the values in $\mathbf{p}_i$, but can only ever decrease the size of the $\mathbf{p}_i$ vector. In particular, due to limited misreports we have $\hat{a}_i \geq a_i$, $\hat{d}_i \leq d_i$ and $\hat{r}_i \leq r_i$, and thus $\mathbf{p}_i^{\langle \hat{a}_i, \hat{d}_i, \hat{r}_i \rangle}$ contains a *subset* of the elements from $\mathbf{p}_i^{\langle a_i, d_i, r_i \rangle}$. As these vectors are by definition increasingly ordered, it follows that $p_{i,j}^{\langle \hat{a}_i, \hat{d}_i, \hat{r}_i \rangle} \geq p_{i,j}^{\langle a_i, d_i, r_i \rangle}, \forall j \leq (\hat{d}_i - \hat{a}_i + 1)$. Since the payment consists of the first $k_i = \hat{k}_i$ elements, this can only increase by misreporting.

$\boldsymbol{\pi_i(\hat{\theta}_i) \neq \pi_i(\theta_i)}$: First, we show that $\pi_i(\hat{\theta}_i) > \pi_i(\theta_i)$ could never occur. Since the threshold values remain the same, but the agent can win fewer units per time step (when reporting a lower maximum consumption rate), and/or the number of time steps in which allocations occur decreases (when reporting a later arrival and/or earlier deadline), Stage 1 of the mechanism can only allocate fewer or equal numbers of units. Furthermore, since $p_{i,j}^{\langle \hat{a}_i, \hat{d}_i, \hat{r}_i \rangle} \geq p_{i,j}^{\langle a_i, d_i, r_i \rangle}$, the possibility of cancelling can only increase in Stage 2. Thus, it always holds that $\pi_i(\hat{\theta}_i) \leq \pi_i(\theta_i)$.

Now, we consider the case $\pi_i(\hat{\theta}_i) < \pi_i(\theta_i)$. First, as shown for the case $\pi_i(\hat{\theta}_i) = \pi_i(\theta_i)$ above, we know that $\sum_{j=1}^{\pi_i(\hat{\theta}_i)} p_{i,j}^{\langle a_i, d_i, r_i \rangle} \leq \sum_{j=1}^{\pi_i(\hat{\theta}_i)} p_{i,k}^{\langle \hat{a}_i, \hat{d}_i, \hat{r}_i \rangle}$ (i.e., the payment for those units won can only increase by misreporting arrival and/or departure). Furthermore, we know that the allocation $\pi_i(\theta_i)$ is preferable to any other allocation $\pi_i(\hat{\theta}_i) < \pi_i(\theta_i)$, otherwise reporting the true valuation vector $\mathbf{v}_i$ would not be a dominant strategy. Since the payment for these items is potentially even higher when misreporting, the agent cannot benefit by winning fewer items. □

Theorem 2. *Given non-increasing marginal valuations and limited misreports, both the on-departure cancellation and immediate cancellation policies with payment policy according to Equation 4 are DSIC.*

Proof. The proof of this theorem follows directly from the above lemmas. Lemmas 1 and 2 show that, for any triple of arrival/departure/consumption rate (mis)-reports, $\langle \hat{a}_i \hat{d}_i, \hat{r}_i \rangle$, the allocation policy is truthful in terms of the valuation vector $\mathbf{v}_i$, given an appropriate payment policy. Furthermore, the payments in Equation 4 correspond to those in Lemma 2, and therefore they truthfully implement the mechanism. Finally, Lemma 3 completes this reasoning, by showing that, for a truthful report of valuation vector $\mathbf{v}_i$, agents cannot benefit from misreporting arrivals/departures. □

4.5 Implications for the PHEV Domain

Since we assume the units to be perishable, it may not always be possible to cancel units once they are allocated. Whereas this is not a problem for immediate cancellation, since units are never allocated to begin with, the on-departure cancellation policy requires the battery to be partially discharged before departure of the vehicle. Although this may be undesirable, it is in the agent's best interest to avoid paying for these units given the design of the mechanism, since the marginal value for these units is less than the marginal payment. An agent could avoid discharging by unplugging before the units are discharged, but then the agent will end up paying for these relatively expensive units. This noted, in Section 7, in addition to comparing the `IM` and `OD` mechanisms, we also evaluate to what extent the mechanisms would be manipulable if they were designed with the simple greedy allocation policy, and without assuming any cancellation.

5. Theoretical Bounds on Allocative Efficiency

An important question given the online nature of the allocation is how the allocative efficiency compares to that of an optimal *offline* allocation, assuming full knowledge of the future. As discussed in Sections 3 – 4.2, in our online setting, it is not possible to achieve an optimal allocation, because agents arrive and leave the market continuously. Moreover, in a multi-dimensional online setting, the allocation of some units needs to be cancelled in order to maintain truthfulness. Nevertheless, the optimal offline allocation represents a useful upper bound of what could be achieved in terms of allocative efficiency, if the preferences and availability constraints of all agents were known in advance.

To this end, in the following, we study the theoretical worst-case performance of both `IM` and `OD`. More precisely, for each of these policies, we consider two types of inefficiencies:

- *Worst-case cancellation ratio.* This is the fraction of units from those allocated to *any single agent* that need to be cancelled in the worst case (and, to maintain the incentive properties, cannot be allocated to any other agent). Formally, let I be the set of all agents present in the market at any time point during the time interval over which the bound is computed. Denote by $\pi_i^{pre}(\theta_I)$ the total number of units pre-allocated to agent $i \in I$ in Stage 1 of the policy over the entire active period of this agent, and by $\pi_i^{canc}(\theta_I) = \pi_i^{pre}(\theta_I) - \pi_i(\theta_I)$ the number of units that were cancelled in Stage 2.

Given this, the cancellation ratio for a specific agent i is $R_{C,i}(\theta_I) = \frac{\pi_i^{canc}(\theta_I)}{\pi_i^{pre}(\theta_I)}$. Then we define the worst-case cancellation ratio over all agents $i \in I$ and types $\theta_I \in \Theta^{|I|}$ as:

$$R_C^{max} = \max_{\theta_I \in \Theta^{|I|}} \max_{i \in I} \frac{\pi_i^{canc}(\theta_I)}{\pi_i^{pre}(\theta_I)}.$$

- *Competitive ratio of allocative efficiency.* Whereas the cancellation ratio considers the worst-case for an individual agent, the competitive ratio compares the *social welfare* of our mechanism with the social welfare achieved by the optimal offline mechanism, with full information about future arrivals. Here, social welfare is defined as the sum of the valuations obtained by all agents (i.e., the sum of utilities excluding any payments). In more detail, following the work of Parkes (2007), the competitive ratio for our setting is defined as follows. Let $\pi_i^{on}(\theta_I)$ denote the number of units allocated by our online mechanism on departure of agent i given the types of all agents, $\theta_I \in \Theta^{|I|}$, and $\pi_i^{off}(\theta_I)$ denote the number of units allocated by the optimal offline mechanism to agent i. The social welfare of the allocations is then defined as: $V^{on}(\theta_I) = \sum_{i \in I} \sum_{k=1}^{\pi_i^{on}} v_{i,k}$ for the online case, respectively $V^{off}(\theta_I) = \sum_{i \in I} \sum_{k=1}^{\pi_i^{off}} v_{i,k}$ for the offline case. Now, a competitive analysis assumes the existence of an "adversary" that can choose from a set of inputs, and in our case the adversary can choose any set of agent types $\theta_I \in \Theta^{|I|}$. Given this, an online mechanism is said to be c-competitive for efficiency, if there exists a constant $c \geq 1$ such that: $\forall \theta_I \in \Theta^{|I|} : \frac{V^{on}(\theta_I)}{V^{off}(\theta_I)} \geq \frac{1}{c}$. We can also say that our online mechanism is guaranteed to achieve within a fraction of $\frac{1}{c}$ of the value of the optimal offline algorithm.

Our motivation for studying these two metrics is as follows. First, as outlined in Section 4.2, both variants of the mechanism we propose require that part of the allocation of some agents is sometimes cancelled, in order to ensure truthfulness. It is natural to ask what is the worst-case fraction of the number of units allocated to any agent that will need to be cancelled, under both types of mechanisms (i.e., with immediate and on-departure cancellation), for any market set-up. For the second criteria (i.e., the competitive ratio of allocative efficiency), we follow the metric proposed by Hajiaghayi et al. (2005) and Parkes (2007) for online domains, with the caveat that deriving this bound for the multi-dimensional case is considerably more involved than for the single-dimensional one, due to the required cancellations. In Section 5.1 we study these issues for the mechanism with immediate cancellation, and in Section 5.2 for the mechanism with on-departure cancellation.

5.1 Worst-Case Bounds for the Mechanism with Immediate Cancellation

The following theorem shows that, when using the online mechanism with immediate cancellation, the worst-case cancellation ratio goes to 1 as the number of units required by a single agent goes to infinity.

Theorem 3. *Using the* `IM` *allocation policy,* $\lim_{n \to \infty} R_C^{max} = 1$*, where* n *is the maximum demand.*

Proof. The proof is by example. Consider the following setting consisting of an agent A with marginal valuation vector $\mathbf{v}_A = \langle v_1, v_2, ... v_n \rangle$, where these values are *strictly* decreasing, i.e., $v_1 > v_2 > .... > v_n$. We assume that this agent arrives at time $a_A = 1$, departs at d_A, where $d_A = n \cdot (n+1)/2$, and has a maximum charging speed of $r_A = 1$. Agent A faces a sequence of cursory (i.e. local) agents, where each of these agents desires exactly one unit, is present in the market for only one timestep and departs immediately afterwards. At any one time there is exactly one of these cursory agents in the market. The valuations of these agents are as follows. The first agent has a valuation of $\mathbf{v}_1 = \langle v_1 - \epsilon \rangle$, the next *two* agents, $i = 2$ and $i = 3$, have valuations of $\mathbf{v}_i = \langle v_2 - \epsilon \rangle$, the next *three* agents, $i \in \{4, 5, 6\}$, have valuations $\mathbf{v}_i = \langle v_3 - \epsilon \rangle$, the next *four* agents, $i \in \{7, 8, 9, 10\}$, $\mathbf{v}_i = \langle v_4 - \epsilon \rangle$, etc. Thus, in total, there is a sequence of $n \cdot (n+1)/2$ of these agents. Here, ϵ is sufficiently small such that $v_1 - \epsilon > v_2$. As a result, agent A imposes the following externality at each timestep: $\langle v_1 - \epsilon, v_2 - \epsilon, v_2 - \epsilon, v_3 - \epsilon, v_3 - \epsilon, v_3 - \epsilon, v_4 - \epsilon, v_4 - \epsilon, v_4 - \epsilon, v_4 - \epsilon, ... \rangle$ (noting that, since the cursory agents are only present in the market for a single time step, cancelling an allocation does not affect the externality, nor the marginal payment).

The allocation for this settings then proceeds as follows. In the first time step, the unit is pre-allocated to agent A (since $v_1 > v_1 - \epsilon$) and there is no cancellation. In the second time step, the unit is again pre-allocated to agent A (since $v_2 > v_2 - \epsilon$), but at this point the marginal payment is $p_i^{\langle 2 \rangle} = \langle v_2 - \epsilon, v_1 - \epsilon \rangle$. Since the marginal value for the second unit is less than the marginal payment for this unit, i.e., $v_2 < v_1 - \epsilon$, the unit gets cancelled. Therefore, at time $t = 3$, the marginal value of agent A is still v_2, and the third unit also gets allocated to the agent, and this time it is not cancelled. However, in the next two time-steps, the units are pre-allocated and cancelled both times. To see this, note that the marginal payment at time $t = 5$ is $p_i^{\langle 2 \rangle} = \langle v_3 - \epsilon, v_3 - \epsilon, v_2 - \epsilon, v_2 - \epsilon, v_1 - \epsilon \rangle$. Since $v_3 < p_{i,3}^{\langle 2 \rangle} = v_2 - \epsilon$, this unit gets cancelled.

More generally, for every k^{th} unit which is allocated and not cancelled, the marginal value of agent A becomes v_{k+1}, and the next k units will be first pre-allocated (since the marginal value of the cursory agents are $v_{k+1} - \epsilon$), but then subsequently cancelled (since these all have a marginal payment of $v_k - \epsilon$). Only the $(k+1)^{th}$ unit will be allocated and not cancelled, but then the next $k+1$ units will be cancelled, and so on.

As a result, all $\pi_A^{pre} = 1 + 2 + 3 + 4 + \ldots + n = n \cdot (n+1)/2$ units will be pre-allocated. Of those units $\pi_A^{canc} = 0 + 1 + 2 + 3 + \ldots + (n-1) = (n-1) \cdot n/2$ will be cancelled, and $\pi_A = n$ will remain allocated. Therefore, the ratio of number of units cancelled as $n \to \infty$ is:

$$R_{C,A} = \lim_{n \to \infty} \frac{\pi_A^{canc}}{\pi_A^{pre}} = \lim_{n \to \infty} \frac{n^2 - n}{n^2 + n} = \lim_{n \to \infty} \frac{n-1}{n+1} = 1$$

□

The above theorem shows that the worst-case result for an individual agent is unbounded. We can use this result (and the example constructed in the proof) to derive a similarly negative result for the allocative efficiency (i.e., the overall efficiency of the system):

Theorem 4. *For the mechanism with immediate cancellation, the competitive ratio of the allocation efficiency is unbounded. That is, there exists no finite c, such that:*

$$\forall \theta_I \in \Theta^{|I|} : \frac{V^{on}(\theta_I)}{V^{off}(\theta_I)} \geq \frac{1}{c}$$

Proof. Generally, there are two potential sources of inefficiency w.r.t. the offline allocation: either the units are pre-allocated and subsequently cancelled, or some units are allocated to agents that have less utility for them than the agents who would be allocated in the offline case. Our proof is based on the former source of inefficiency and uses the same example given in Theorem 3.

In the example of Theorem 3 we showed that it is possible to construct an example where, given $n \cdot (n+1)/2$ units of supply, $(n-1) \cdot n$ of these are cancelled (and thus not allocated to any agent) and n are allocated. Now, suppose that the valuations for these units by all agents (including agent A) are between $[\underline{v}, \overline{v}]$, where $\overline{v}/\underline{v} = r$ is a finite constant. Since using the optimal offline allocation all units will be allocated, the total value will be *at least*: $V^{off} \geq \underline{v} \cdot n \cdot (n+1)/2$. On the other hand, the online allocation using the online allocation using immediate cancellation will have a value of *at most*: $V^{on} \leq \overline{v} \cdot n$.

Given this, the following holds:

$$\lim_{n\to\infty} \frac{V^{on}}{V^{off}} \leq \lim_{n\to\infty} \frac{\overline{v} \cdot n}{\underline{v} \cdot n \cdot (n+1)/2} = \lim_{n\to\infty} \frac{2r}{n+1} = 0$$

Therefore, for any constant c, it is always possible to find a counter example where the worst-case efficiency is lower than $1/c$. □

Thus, there is no theoretical bound on the efficiency loss when using the immediate cancellation allocation policy. However, the proof relies on there being an agent who is infinitely patient, has infinite demand, and has a higher valuation than all other bidders for each unit. In practice, such an extreme situation would never occur. To consider more practical scenarios, therefore, in Section 7 we use simulations to investigate realistic settings. After showing the worst-case bounds with immediate cancellation, in the remainder of Section 5 we derive theoretical bounds for the on-departure cancellation mechanism. Specifically, we will show that this mechanism provides much better bounds. In fact, the competitive bounds for the efficiency are the same as those for single-unit demand settings, where no cancellation occurs.

5.2 Worst-Case Bounds for the Mechanism with On-Departure Cancellation

This section is divided into two parts: in Section 5.2.1 we discuss the worst-case cancellation ratio for a particular agent and provide a tight bound, while in Section 5.2.2 we consider the bound on allocative efficiency for the entire market.

5.2.1 Worst-Case Cancellation Ratio

This section is organised as follows. First, we show that at most half of the units are cancelled for any particular agent. Then we go on to show that there exist examples where half of them are cancelled. Note that, for convenience, the following lemma is formulated in terms of units retained instead of units cancelled.

Lemma 4. *Using on-departure cancellation, suppose that an agent i is pre-allocated n units by the departure time d_i, and k of these units are kept in Stage 2 (and so the mechanism cancels $n-k$ units). Then, for any type profile $\theta_I \in \Theta^{|I|}$, and any agent $i \in I$, $k \geq \lceil n/2 \rceil$ (i.e., at least half of the units are allocated).*

Proof. To prove this property, we start by deriving two inequalities which hold for any value of k. First, since k is defined as the number of units kept, and the remaining ones $(n-k)$ are cancelled, it must hold that $v_{i,k+1} < p_{i,k+1}$ (otherwise the $(k+1)^{th}$ unit would not be cancelled, contradicting the definition).

The second inequality is given by $v_{i,k+1} \geq p_{i,n-k}$ and is less obvious. To see why this always holds, we need an observation about how greedy allocation works. Recall that n is the number of units pre-allocated by the greedy allocation policy. Therefore, each of the active marginal values, $v_{i,1}, \ldots, v_{i,n}$, were at some point $t \in [a_i, d_i]$ among the top $S(t)$ highest marginal values. Consequently, the lowest marginal value of the ones pre-allocated, $v_{i,n}$, must be greater than the marginal payment of at least one unit (otherwise it could not have won the unit). Since the marginal payments are sorted in an increasing order, it must therefore hold that $v_{i,n} \geq p_{i,1}$. Similarly, for the next-lowest value, it must hold that $v_{i,n-1} \geq p_{i,2}$, and so on. In general we can write $v_{i,n-j+1} \geq p_{i,j}$, $j \in \{1, n\}$. If we set $j = n-k$, we get $v_{i,k+1} \geq p_{i,n-k}$.

Therefore, in order for the greedy policy to allocate n units and for the mechanism to subsequently cancel the units from positions $k+1$ to n (assuming $k+1 < n$, otherwise no cancelling will take place on departure of agent i), the following inequalities must be satisfied:

$$\begin{cases} v_{i,k+1} < p_{i,k+1} \\ v_{i,k+1} \geq p_{i,n-k} \end{cases} \tag{5}$$

Given this, we now show that $k \geq \lceil n/2 \rceil$ by contradiction. Suppose that $k = \lceil n/2 \rceil - 1$, i.e., strictly more than $\lceil n/2 \rceil$ are cancelled. Then the above conditions become:

$$\begin{cases} v_{i,\lceil n/2 \rceil} < p_{i,\lceil n/2 \rceil} \\ v_{i,\lceil n/2 \rceil} \geq p_{i,n-\lceil n/2 \rceil+1} \end{cases} \tag{6}$$

To show the contradiction, we need to consider separately the cases where n is even and where n is odd. If n is even, then we have $n = \lceil n/2 \rceil + \lceil n/2 \rceil$, and the above system becomes:

$$\begin{cases} v_{i,\lceil n/2 \rceil} < p_{i,\lceil n/2 \rceil} \\ v_{i,\lceil n/2 \rceil} \geq p_{i,\lceil n/2 \rceil+1} \end{cases} \tag{7}$$

This implies $p_{i,\lceil n/2 \rceil+1} < p_{i,\lceil n/2 \rceil}$, but since the marginal price vector $\mathbf{p}_i$ is weakly increasing by definition, this leads to a contradiction. For the case where n is odd, we have that $n = \lceil n/2 \rceil + \lceil n/2 \rceil - 1$, and the conditions become:

$$\begin{cases} v_{i,\lceil n/2 \rceil} < p_{i,\lceil n/2 \rceil} \\ v_{i,\lceil n/2 \rceil} \geq p_{i,\lceil n/2 \rceil} \end{cases} \tag{8}$$

Clearly, both equations cannot be satisfied simultaneously, leading to a contradiction.

Note that any value of $k < \lceil n/2 \rceil$ would lead to such a contradiction due to p_i being increasing, hence we necessarily have that $k \geq \lceil n/2 \rceil$, completing the proof. □

To complete our analysis, we show that this bound is tight, i.e., that there exist settings in which half of the units allocated to an agent are cancelled.

Lemma 5. *There exist settings in which the mechanism with on-departure cancellation cancels the allocation of $\lceil n/2 \rceil$ units, or no more than $\lceil n/2 \rceil$ units are kept, on departure of an agent i, where n is the number of pre-allocated units.*

Proof. The proof for this is done by constructing such a worst-case example. Consider a single agent, A, who is in the market for n time periods, and has a demand for n units, where n is even. The first $n/2$ marginal valuations are equal to 4, and the remaining ones are 2. For example, for $n = 8$, the marginal valuation vector becomes $\mathbf{v}_A = \langle 4, 4, 4, 4, 2, 2, 2, 2 \rangle$. Similar to the proof of Theorem 3, this agent is faced, in each time step, by a different, single cursory agent, which participates only in that time step. The valuations of the first $n/2$ cursory agents in the sequence is given by $\mathbf{v}_i = \langle 3 \rangle$, and the second half of the agents has $\mathbf{v}_i = \langle 1 \rangle$. Thus, for $n = 8$, the marginal payment of agent A would be $\mathbf{p}_{-A} = \langle 1, 1, 1, 1, 3, 3, 3, 3 \rangle$. In such a setting, the mechanism with on-departure cancellation would pre-allocate all units to agent A (since $v_{A,k} = 4 > v_{k,1} = 3$ for $k \leq n/2$ and $v_{A,k} = 2 > v_{k,1} = 1$ for $k > n/2$). However, on departure of agent A, exactly half of the units allocated are cancelled (since $v_{A,k} = 2 < \mathbf{p}_{-A,k} = 3$ for $k > n/2$). □

Finally, we unify the results from Lemmas 4 and 5 in the following theorem.

Theorem 5. *In a setting with on-departure cancellation and non-increasing marginal values, for any number of units and agents present, the worst case cancellation-ratio for the number of units allocated to an agent i is $R_C^{max} = \frac{1}{2}$.*

Proof. Lemma 4 shows that, regardless of set-up, no more than half of the units allocated to any agent can be cancelled on its departure, thus $R_C^{max} \leq \frac{1}{2}$, regardless of the setting (i.e., the possible input types of the agents). Lemma 5 shows there exist settings where the cancellation ratio is exactly $R_C^{max} = \frac{1}{2}$, completing the proof. □

Note that in practice and for smaller settings, significantly fewer than half of the units are cancelled. The worst case cancellation ratio 1/2 allocations occurs only in a very specifically constructed example, and, as shown in the experimental analysis, for most realistic distributions in our application domain, the actual performance is much better.

5.2.2 Competitive Bound on Allocative Efficiency

The previous section discusses the cancellation problem from the perspective of single agents, not the whole market. In this section, we show that, in the case that agents have weakly decreasing marginal values, the allocation returned by the on-departure cancellation mechanism is 2-competitive with the optimal offline allocation. This result means that the multi-unit demand case with on-departure cancellation is no worse in terms of worst-case competitive bound than the single-unit demand problem discussed by Hajiaghayi et al. (2005), and Parkes (2007), despite the fact that for a single unit demand there is no need for cancellation to ensure incentive compatibility. Formally, we can state this through the following theorem:

Theorem 6. *The mechanism with on-departure cancellation is 2-competitive with the optimal offline allocation, for a setting with non-increasing marginal values.*

Proof. In order to establish a competitive bound with the optimal offline allocation, we use a "charging argument" similar to that of Hajiaghayi et al. (2005).[5] The basic idea is to "charge" (or match) all the marginal value units of each agent that are allocated in the offline case with another, higher-valued unit that is allocated both offline and online. This is either the unit itself, or the higher value unit that causes it not to be allocated in the online market. Formally, consider all units $v_{i,p}$ (belonging to some agent i in position p) which are allocated *both* in the offline and online case. If each such unit $v_{i,p}$ can be charged at most twice, once to itself, and once to a lower valued unit allocated offline but not online, then it follows that the worst-case social welfare ratio between the the online vs. offline allocation cannot drop below 1:2.

Now, for agents with single-unit demand (such as the case discussed by Hajiaghayi et al., 2005), it is easy to see this property always holds, because each unit $v_{i,p}$ can be allocated online at most once, thus it can displace *at most* one other unit $v_{j,q}$. Crucially, no units are cancelled. In a multi-unit demand setting, the argument becomes more involved, because each unit $v_{i,p}$ (allocated both online and offline) can affect the online market in several ways:

- It can displace another unit $v_{j,q}$ that would be allocated offline, where by "displace" we mean specifically that unit $v_{j,q}$ is never *pre-allocated* online (hence cancellation does not apply to it).

- It can cause the cancellation of another unit $v_{j,q}$. In this second case, unit $v_{j,q}$ is pre-allocated, but its allocation is cancelled due to the presence of unit $v_{i,p}$ in the market (meaning that its pre-allocation would not have been cancelled on departure of agent j, if unit $v_{i,p}$ were not present).

The main issue that remains to be shown is that unit $v_{i,p}$ can only displace or cause the cancellation of *at most* one other unit that would be allocated offline. Thus, it cannot displace two or more other units that are allocated offline, but are not allocated online, due to the presence of unit $v_{i,p}$.

We show this by contradiction. Formally, suppose there are three units: $v_{i,p}$, $v_{j,q}$ and $v_{k,r}$ all allocated in the offline case (with $v_{j,q} < v_{i,p}$ and $v_{k,r} < v_{i,p}$). Unit $v_{i,p}$ is allocated in the online case (i.e., pre-allocated and not cancelled). Units $v_{j,q}$ and $v_{k,r}$ are not allocated in the online case if unit $v_{i,p}$ is present, but are allocated online if unit $v_{i,p}$ is not present. Given this set-up, there are three possible cases:

1. Neither units $v_{j,q}$ or $v_{k,r}$ are pre-allocated online when unit $v_{i,p}$ is present (hence, there is no cancellation of either $v_{j,q}$ or $v_{k,r}$).

2. Unit $v_{j,q}$ is never pre-allocated online, but unit $v_{k,r}$ is pre-allocated and its allocation is cancelled later (i.e., on departure of agent k from the market), if unit $v_{i,p}$ is present.

3. Both units $v_{j,q}$ and $v_{k,r}$ are pre-allocated, but their pre-allocations are cancelled on the departure from the market agents j, respectively k, if unit $v_{i,p}$ is present.

5. Here the term "charging" does not refer to electricity charging, but represents the name of a proof device used in online mechanism design.

In all cases, if unit $v_{i,p}$ is not present, both units $v_{j,q}$ and $v_{k,r}$ are pre-allocated and not cancelled in the online case. In order to complete the proof we need to show, by contradiction, that each of these three cases could not occur.

Case 1 is very similar to the case of single unit demand discussed by Hajiaghayi et al. (2005), as no cancellation occurs for these units. It is relatively straightforward to see this cannot occur, as any unit $v_{i,p}$ can be pre-allocated at most once (at some time t), thus it can displace at most one other unit that would have been allocated otherwise. This can be either the unit allocated online at time time t, if this unit is allocated online later on, the unit which is, in turn, displaced by it.

Case 2: Suppose that $v_{j,q}$ (belonging to some agent j active between $[a_j, d_j]$) is the unit assumed not pre-allocated at all when $v_{i,p}$ is present, and unit $v_{k,r}$ is the unit that is allocated and then cancelled. There are two subcases to consider here, which require separate discussion.

Case 2A: First, consider that $v_{j,q} > v_{k,r}$. In this case, agent k has a lower marginal value than that of agent j, but its value $v_{k,r}$ is still pre-allocated by our essentially greedy allocation policy, while $v_{j,q}$ is not. This means that agent k must be more patient than agent j, hence $d_j < d_k$, otherwise $v_{j,q}$ would have been pre-allocated instead.

Now, if we denote by $\mathbf{p}_k$ the payment vector of agent k, defined as in Section 4.2. For unit $v_{k,r}$ to be cancelled it must hold that $v_{k,r} < p_{k,r}$. Now, denote by $\mathbf{p}_k^{<-i>}$ the vector of marginal payments of agent k when agent i is not present in the market, and recall our assumption that now the value $v_{j,q}$ is allocated. Thus, we have:

$$\mathbf{p}_k^{<-i>} = \mathrm{incr}(\mathbf{p}_k \backslash \{v_{i,p}\} \cup \{v_{j,q}\})$$

where incr is the operator that orders elements in increasing order. Since $v_{j,q} > v_{k,r}$, it follows that $v_{k,r}^{<-i>} < p_{k,r}$, thus the allocation of unit $v_{k,r}$ would still be cancelled, even without unit $v_{i,p}$.

Case 2B: For the second subcase, we consider $v_{j,q} \leq v_{k,r}$, i.e., the value of the unit that is displaced by agent i is lower than that the one pre-allocated and cancelled. First note that, for this case to occur, unit $v_{j,q}$ but be allocated online within $[a_k, d_k]$, the active window of agent k. This is an obvious condition: if agent j is allocated online outside this window (and is displaced there when agent i is in the market, but the displacement occurs outside $[a_k, d_k]$), then units $v_{j,q}$ or $v_{i,p}$ cannot influence the cancellation of unit $v_{k,r}$ (because each unit can be pre-allocated at most once, and in this case the pre-allocation of $v_{i,p}$ would happen outside $[a_k, d_k]$).

As previously, recall the condition for unit $v_{k,r} < p_{k,r}$, required for unit $v_{k,r}$ to be cancelled. Note that this means there are at least k units between $[a_k, d_k]$ that are higher in value than unit $v_{k,r}$, and thus, in an *offline allocation* (which is our benchmark) would need to take priority over it. In our setting, one of these units is $v_{i,p}$. But even after removing $v_{i,p}$ from it, vector $\mathbf{p}_k^{<-i>} = \mathrm{incr}(\mathbf{p}_k \backslash \{v_{i,p}\} \cup \{v_{j,q}\})$ must contain at least $k-1$ values higher than $v_{k,r}$. In an offline allocation *without* unit $v_{i,p}$, these $k-1$ values must be given priority, together with at least unit $v_{k,r}$. However, this means that unit $v_{j,q}$ cannot be allocated *offline* between $[a_k, d_k]$, being lower in value than $v_{k,r}$. This gives a contradiction with our initial assumption that *both* units $v_{i,j}$ and $v_{k,r}$ are allocated *offline* (as well as online) without unit $v_{i,p}$ present.

To explain this intuitively, what this means is that a unit $v_{i,p}$ can cause both the displacement (non-allocation) of a unit $v_{j,q}$ and the cancellation of another one $v_{k,r}$, but it's not possible that both of these units were high enough value to be allocated in the offline case as well. Thus, at most one *other* offline-allocated unit is not allocated online because of the presence of unit $v_{i,p}$ in the market.

Case 3: In this final case, both units would need to be pre-allocated and cancelled, in the absence of value $v_{i,p}$. The contradiction for this case can be shown similarly to Case 2A from above. Considering the marginal price vectors of agents j and k without $v_{i,p}$ in the market:

$$\mathbf{p}_j^{<-i>} = \text{incr}(\mathbf{p}_j \backslash \{v_{i,p}\} \cup \{v_{k,r}\})$$

$$\mathbf{p}_k^{<-i>} = \text{incr}(\mathbf{p}_k \backslash \{v_{i,p}\} \cup \{v_{j,q}\})$$

It is easy to see that, regardless whether the value $v_{j,q}$ or $v_{k,r}$ is lower, for that value the cancellation would still occur on departure in a market without agent i, leading to a contradiction.

To summarise, we have now exhaustively shown that the contradiction holds in all possible cases. Thus, a unit $v_{i,p}$ allocated both online and offline can at most displace (or lead to the cancellation of) *one* other unit allocated *offline*. Thus, at most two units allocated offline can be charged to any unit allocated online, completing the proof. □

6. Computational Aspects

In this section, we consider the implications of implementing the mechanisms in practice, including the computational complexity of our algorithms. We will examine both on-departure and immediate cancellation separately, as they differ fundamentally in their complexity.

6.1 Implementing On-Departure Cancellation

Algorithm 1 briefly outlines an implementation of our mechanism with on-departure cancellation (`OD`). Here, we assume that the first time step is denoted by t_0, the second by $t_1 = t_0 + 1$, and so on. For simplicity, we use I throughout this section to denote the full set of all agents arriving over all time points, but we note that no algorithm explicitly uses information about future arrivals. Initially, the algorithm sets the endowments of all agents to 0 (line 2), as no units have been allocated. Then, for every time step t, the algorithm first pre-allocates units using the greedy allocation policy (line 4). This can be done in $O(N \cdot r_{\max})$ using the well known linear-time selection algorithm described by Blum, Floyd, Pratt, Rivest, and Tarjan (1973), where $N = |I|$ is the total number of agents and $r_{\max} = \max_{i \in I} \hat{r}_i$ is the maximum consumption rate.

Next, the algorithm computes the marginal payments up to time step t by rerunning the market without each active agent (line 6). In rerunning the market for a particular agent i, it is important to note that only the pre-allocations have an effect on $\mathbf{p}_i^{\langle t \rangle}$ — any cancellations are irrelevant, because they do not affect the future development of the market. Therefore, it is only necessary to compute the greedy allocation with each agent i removed, which has a total run-time, for all active agents, in $O(N^2 \cdot r_{\max})$ (assuming results from previous time steps are re-used). Updating $\mathbf{p}_i^{\langle t \rangle}$ with the new marginal payments can be

Algorithm 1 Mechanism with On-Departure Cancellation (`OD`).

1: **procedure** OnDepartureMechanism(θ_I, S)
2: $\quad \mathbf{k}^{\langle t_0 \rangle} \leftarrow \langle 0, 0, \ldots, 0 \rangle$ $\quad \triangleright$ Initial endowments, $k_i^{\langle t_0 \rangle} = 0$, for all $i \in I$
3: $\quad$ **for all** $t \in \{t_0, t_1, \ldots\}$ **do**
4: $\quad\quad \mathbf{k}^{\langle t+1 \rangle} \leftarrow$ GreedyAllocation($\theta_I, S(t), \mathbf{k}^{\langle t \rangle}$) $\quad \triangleright$ Run greedy allocation
5: $\quad\quad$ **for all** $i \in \{j \in I | \hat{a}_j \leq t \wedge \hat{d}_j \geq t\}$ **do** $\quad \triangleright$ Iterate through active agents
6: $\quad\quad\quad$ update $\mathbf{p}_i^{\langle t \rangle}$ using θ_{-i} $\quad \triangleright$ Run market without i
7: $\quad\quad\quad$ **if** $\hat{d}_i = t$ **then** $\quad \triangleright$ If agent is departing
8: $\quad\quad\quad\quad k_i, p_i \leftarrow k_i^{\langle t+1 \rangle}, p_i^{\langle t \rangle}$ $\quad \triangleright$ Final pre-allocation and marginal payments
9: $\quad\quad\quad\quad$ **while** $v_{i,k_i} < p_{i,k_i}$ **do**
10: $\quad\quad\quad\quad\quad k_i \leftarrow k_i - 1$ $\quad \triangleright$ Cancel units while necessary
11: $\quad\quad\quad\quad$ **end while**
12: $\quad\quad\quad\quad x_i(\theta_{-i} | k_i) \leftarrow \sum_{k=1}^{k_i} p_{i,k}$ $\quad \triangleright$ Final payment for agent i
13: $\quad\quad\quad$ **end if**
14: $\quad\quad$ **end for**
15: $\quad$ **end for**
16: **end procedure**

done using a simple insertion algorithm, and, noting that only the lowest $|\mathbf{v}_i|$ payments for each agent i need to be kept, this can be done, for all agents, in $O(N \cdot r_{\max} \cdot v_{\max})$, where $v_{\max}$ is the maximum length of any agent's valuation vector.

Finally, on departure, in lines 8–11, any units that have a lower valuation than their corresponding marginal payments are cancelled, and the final payment is calculated in line 12. This can be done in $O(N \cdot v_{\max})$ by simply iterating through the values.

In summary, the time complexity of our algorithm for the `OD` mechanism is $O(N^2 \cdot r_{\max} + N \cdot r_{\max} \cdot v_{\max} + N \cdot v_{\max})$ for each time step. If $r_{\max}$ and $v_{\max}$ are assumed to be constant[6], this simplifies to $O(N^2)$. Generally, this means that the algorithm can be executed quickly, even for large numbers of agents.

6.2 Implementing Immediate Cancellation

Next, we consider our mechanism with immediate cancellation (`IM`), as shown in Algorithm 2. The key difference to the `OD` mechanism here is that units are potentially cancelled at every time step that an agent is active (lines 7–9), rather than only on departure. This small modification has a significant impact on the computational tractability of the mechanism. Unlike the previous mechanism, when computing the marginal payments in line 6, the cancellations of other agents now affect the payments. This feature was already highlighted in the example in Section 4.3.1, where the cancellation of the second unit pre-allocated to agent 1 causes a change in the marginal payments to agent 3. More generally, cancellations have an immediate effect on the endowments of agents and this directly affects the active marginal valuations $\mathcal{V}^{\langle t \rangle}$ in subsequent time steps.

6. This is reasonable, as they are limited by technological constraints in practice. In particular, $r_{\max}$ is limited by battery and infrastructure constraints, while $v_{\max}$ is related to the petrol savings achievable by an EV (as will be detailed in Section 7.3.3).

Algorithm 2 Mechanism with Immediate Cancellation (`IM`).

1: **procedure** ImmediateMechanism(θ_I, S)
2: $\quad \mathbf{k}^{\langle t_0 \rangle} \leftarrow \langle 0, 0, \ldots, 0 \rangle$ $\triangleright$ Initial endowments, $k_i^{\langle t_0 \rangle} = 0$, for all $i \in I$
3: $\quad$ **for all** $t \in \{t_0, t_1, \ldots\}$ **do**
4: $\quad\quad \mathbf{k}^{\langle t+1 \rangle} \leftarrow$ GreedyAllocation($\theta_I, S(t), \mathbf{k}^{\langle t \rangle}$) $\triangleright$ Run greedy allocation
5: $\quad\quad$ **for all** $i \in \{j \in I | \hat{a}_j \leq t \wedge \hat{d}_j \geq t\}$ **do** $\triangleright$ Iterate through active agents
6: $\quad\quad\quad$ update $\mathbf{p}_i^{\langle t \rangle}$ using θ_{-i} $\triangleright$ Run market without i
7: $\quad\quad\quad$ **while** $v_{i,k_i^{\langle t+1 \rangle}} < p^{\langle t \rangle}_{i,k_i^{\langle t+1 \rangle}}$ **do**
8: $\quad\quad\quad\quad k_i^{\langle t+1 \rangle} \leftarrow k_i^{\langle t+1 \rangle} - 1$ $\triangleright$ Cancel units immediately
9: $\quad\quad\quad$ **end while**
10: $\quad\quad\quad$ **if** $\hat{d}_i = t$ **then** $\triangleright$ If agent is departing
11: $\quad\quad\quad\quad k_i, p_i \leftarrow k_i^{\langle t+1 \rangle}, p_i^{\langle t \rangle}$ $\triangleright$ Final allocation and marginal payments
12: $\quad\quad\quad\quad x_i(\theta_{-i}|k_i) \leftarrow \sum_{k=1}^{k_i} p_{i,k}$ $\triangleright$ Final payment for agent i
13: $\quad\quad\quad$ **end if**
14: $\quad\quad$ **end for**
15: $\quad$ **end for**
16: **end procedure**

Now, in order to determine cancellations when rerunning the market without each active agent i, it is necessary to again compute the marginal payments for all agents in those markets (effectively executing the full algorithm again with θ_{-i}). Clearly, this leads to a recursion that potentially sees all possible subsets of agents evaluated. In the worst case, therefore, the cancellation decisions need to be executed for every agent in every possible subset of I, or $N \cdot 2^{(N-1)}$ times.[7] Simplifying and again assuming $v_{\max}$ and $r_{\max}$ to be constant, this leads to a runtime complexity of $O(N \cdot 2^N)$.

A runtime that is exponential in the number of agents is clearly a problem when applying the mechanism in realistic settings with more than a handful of agents. However, to tackle such problems, it is possible to use a technique akin to branch-and-bound that enters the recursion in line 6 only when necessary. We present this in the following section.

6.3 Speeding Up Immediate Cancellation Using Bounds

To obtain a faster algorithm for the `IM` mechanism, instead of calculating all marginal payments $\mathbf{p}_i^{\langle t \rangle}$ at every time step, we find and iteratively refine lower and upper bounds for these payments. The intuition behind this approach is to choose initial bounds that are easily calculated without resorting to recursion. If an agent's reported valuation for a pre-allocated unit, $v_{i,k}$, lies outside the bounds, we can immediately determine whether the unit is cancelled or not. On the other hand, if $v_{i,k}$ lies between the bounds, we further refine

7. In practice, this recursion only occurs over the set of agents that are active at the same time as the agent i that is being evaluated, as previous decisions are not affected by agent i's presence and similarly agents arriving after $\hat{d}_i$ have no affect on i. This means that settings with large numbers of agents may still be tractable if there is little overlap between the active agents, but for the sake of the analysis in this section, we assume the worst case, that all N agents are active concurrently.

them by iteratively calculating the actual marginal payments for some of the agent's active time steps until the cancellation decision is unambiguous.

In more detail, we use $\mathbf{p}_i^{\downarrow\langle t,s\rangle}$ and $\mathbf{p}_i^{\uparrow\langle t,s\rangle}$ to denote the lower and upper bounds, respectively, for $\mathbf{p}_i^{\langle t\rangle}$. Here, s indicates the level of refinement, or up to which time step the actual marginal payments have been calculated, with $\hat{a}_i - 1 \leq s \leq t$. Analogous to the actual $\mathbf{p}_i^{\langle t\rangle}$ vector, $\mathbf{p}_i^{\downarrow\langle t,s\rangle}$ and $\mathbf{p}_i^{\uparrow\langle t,s\rangle}$ are vectors in increasing order, and they represent bounds for $\mathbf{p}_i^{\langle t\rangle}$, such that $p_{i,k}^{\downarrow\langle t,s\rangle} \leq p_{i,k}^{\langle t\rangle} \leq p_{i,k}^{\uparrow\langle t,s\rangle}$, for all k and s. As the level of refinement, s, is increased, the bounds become tighter, eventually converging to $p_{i,k}^{\langle t\rangle}$. In the following, we describe in detail how to calculate the initial bounds (with $s = \hat{a}_i - 1$, indicating that no actual payments have been calculated yet):

- In order to calculate the initial *lower bounds*, $\mathbf{p}_i^{\downarrow\langle t,\hat{a}_i-1\rangle}$, we rerun the market without i from $\hat{a}_i$ to t using only the greedy allocation policy *without* cancellations. The marginal payments in this market (as described in Section 4.2) are then used as lower bounds. These payments are not necessarily the same as the actual payments $\mathbf{p}_i^{\langle t\rangle}$, as there may be cancellations in the latter, which cause the active marginal valuations to change in subsequent time steps. However, it can be seen easily that they indeed represent a lower bound. Specifically, any cancellations either have no influence or cause an *increase* in one or more elements in the externality vector $\mathcal{E}^{\langle t\rangle}$ (since agents have a lower endowment after cancellations and therefore an equal or higher value for obtaining additional units).

- To calculate the initial *upper bounds*, $\mathbf{p}_i^{\uparrow\langle t,\hat{a}_i-1\rangle}$, we now consider the *actual* market *including* agent i from time step $\hat{a}_i$ to t, calculate the externality that agent i would impose on others for winning each available unit throughout this time interval and then use these to derive upper bounds for the payments. More formally, we define a new multiset set of valuations, $\mathcal{V}_{-i\cup 0}^{\uparrow\langle t\rangle}$, which is derived by simply removing all elements corresponding to agent i from $\mathcal{V}^{\langle t\rangle}$ and padding it with zeroes, if necessary, until its size is at least $S(t)$. Then, we proceed in a similar manner as in Section 4.2 by defining the externality agent i would impose on others in time step t as $\mathcal{E}_i^{\uparrow\langle t\rangle} = \min_{\langle r_i\rangle}(\max_{\langle S(t)\rangle} \mathcal{V}_{-i\cup 0}^{\uparrow\langle t\rangle})$. Given this, the final vector of upper bounds is then $\mathbf{p}_i^{\uparrow\langle t,\hat{a}_i-1\rangle} = \mathrm{incr}(\bigcup_{t'=\hat{a}_i}^{t} \mathcal{E}_i^{\uparrow\langle t'\rangle})$.

  Unlike the actual marginal payments, these upper bounds now include the effect agent i has on the market, as it may be allocated units that would have been allocated to others, or, through its presence, cause the cancellation of other units, which, in turn, affect the active valuations of other agents in subsequent time steps. However, as its only effect is to reduce the supply available to other agents, including agent i in the market can only increase the active valuations of other agents, and therefore $\mathbf{p}_i^{\uparrow\langle t,\hat{a}_i-1\rangle}$ is an upper bound for $\mathbf{p}_i^{\langle t\rangle}$.

Given these bounds, we can now quickly test if an agent's marginal valuation $v_{i,k}$ for a pre-allocated unit falls outside these bounds. If $v_{i,k} < p_{i,k}^{\downarrow\langle t,\hat{a}_i-1\rangle}$, then the unit is cancelled immediately, while if $v_{i,k} \geq p_{i,k}^{\uparrow\langle t,\hat{a}_i-1\rangle}$, it is definitely not cancelled. However, if

Algorithm 3 Mechanism with Immediate Cancellation (`IM`) and Bounds.

procedure BoundedImmediateMechanism(θ_I, S)
  $\mathbf{k}^{\langle t_0\rangle} \leftarrow \langle 0, 0, \ldots, 0\rangle$ ▷ Initial endowments, $k_i^{\langle t_0\rangle} = 0$, for all $i \in I$
  for all $i \in I$ **do**
    $s_i \leftarrow \hat{a}_i - 1$ ▷ Initial refinement of bounds
  end for
  for all $t \in \{t_0, t_1, \ldots\}$ **do**
    $\mathbf{k}^{\langle t+1\rangle} \leftarrow$ GreedyAllocation($\theta_I, S(t), \mathbf{k}^{\langle t\rangle}$) ▷ Run greedy allocation
    for all $i \in \{j \in I | \hat{a}_j \leq t \wedge \hat{d}_j \geq t\}$ **do** ▷ Iterate through active agents
      Calculate $\mathbf{p}_i^{\downarrow\langle t,s_i\rangle}$ and $\mathbf{p}_i^{\uparrow\langle t,s_i\rangle}$ ▷ Add initial bounds for time t
      repeat
        if $v_{i,k_i^{\langle t+1\rangle}} < p^{\downarrow\langle t,s_i\rangle}_{i,k_i^{\langle t+1\rangle}}$ **then**
          $k_i^{\langle t+1\rangle} \leftarrow k_i^{\langle t+1\rangle} - 1$ ▷ Unit definitely cancelled
        else if $v_{i,k_i^{\langle t+1\rangle}} < p^{\uparrow\langle t,s_i\rangle}_{i,k_i^{\langle t+1\rangle}}$ **then**
          $s_i \leftarrow s_i + 1$ ▷ Refine bounds
          Update $\mathbf{p}_i^{\downarrow\langle t,s_i\rangle}$ and $\mathbf{p}_i^{\uparrow\langle t,s_i\rangle}$
        end if
      until $v_{i,k_i^{\langle t+1\rangle}} \geq p^{\uparrow\langle t,s_i\rangle}_{i,k_i^{\langle t+1\rangle}}$
      if $\hat{d}_i = t$ **then** ▷ If agent is departing
        $k_i \leftarrow k_i^{\langle t+1\rangle}$ ▷ Final allocation
        while $\exists x \in \{1, 2, \ldots, k_i\}, p^{\downarrow\langle t,s_i\rangle}_{i,x} \neq p^{\uparrow\langle t,s_i\rangle}_{i,x}$ **do**
          $s_i \leftarrow s_i + 1$ ▷ Refine for final payments
        end while
        $x_i(\theta_{-i}|k_i) \leftarrow \sum_{k=1}^{k_i} p^{\uparrow\langle t,s_i\rangle}_{i,k}$ ▷ Final payment for agent i
      end if
    end for
  end for
end procedure

$p^{\downarrow\langle t,\hat{a}_i-1\rangle}_{i,k} \leq v_{i,k} < p^{\uparrow\langle t,\hat{a}_i-1\rangle}_{i,k}$, then the bounds are ambiguous and need to be further refined. We obtain these refined bounds $\mathbf{p}_i^{\downarrow\langle t,s\rangle}$ and $\mathbf{p}_i^{\uparrow\langle t,s\rangle}$ by computing the actual marginal payments up to some specified time s, after which the bounds are calculated as above. This effectively replaces some of the initial bounds with actual marginal payments, resulting in more accurate overall bounds. Eventually, when $s = t$, we have $\mathbf{p}_i^{\downarrow\langle t,s\rangle} = \mathbf{p}_i^{\langle t\rangle} = \mathbf{p}_i^{\uparrow\langle t,s\rangle}$.

More formally, the refined lower bounds are calculated as $\mathbf{p}_i^{\downarrow\langle t,s\rangle} = \text{incr}((\bigcup_{t'=\hat{a}_i}^{s} \mathcal{E}_i^{\langle t'\rangle}) \cup (\bigcup_{t'=\hat{a}_i+1}^{t} \mathcal{E}_i^{\downarrow\langle t'\rangle}))$, where $\mathcal{E}_i^{\langle t\rangle}$ is the externality vector in the market without i, using immediate cancellations (as used for actual payments), while $\mathcal{E}_i^{\downarrow\langle t\rangle}$ is the corresponding vector in the market without i and without any cancellations. Similarly, the refined upper bounds are calculated as $\mathbf{p}_i^{\uparrow\langle t,s\rangle} = \text{incr}((\bigcup_{t'=\hat{a}_i}^{s} \mathcal{E}_i^{\langle t'\rangle}) \cup (\bigcup_{t'=\hat{a}_i+1}^{t} \mathcal{E}_i^{\uparrow\langle t'\rangle}))$, where $\mathcal{E}_i^{\uparrow\langle t\rangle}$ is as defined above.

Full details for the `IM` mechanism using bounds are given in Algorithm 3. This keeps track of the level of refinement for the bounds of each agent i (as s_i, which is initialised to $\hat{a}_i - 1$ in line 4). Then, instead of updating the actual payments, lines 10–17 repeatedly compare the marginal valuation of the last unit that was pre-allocated ($v_{i,k_i^{\langle t+1 \rangle}}$) to the current upper and lower bounds, refining them as necessary. Note here, that previous calculations of the bounds can be reused, as the algorithm checks them iteratively, either increasing t (when the next time step is calculated) or s_i by 1 (when the bounds are refined). In each case, this means that only the active valuations of one additional time step have to be added to the existing bounds. This loop repeats until the statement in line 17 becomes true, which captures both cases when the last pre-allocated unit is definitely not cancelled (in which case any other pre-allocated units will also not be cancelled, as their marginal valuations will be at least as high, while their respective payments will be equal or lower), or when all pre-allocated units in this time step are cancelled.

Finally, when an agent departs, its final allocation and payments are calculated in lines 18–24. Here, it is now important to calculate the actual marginal payments for all k_i allocated units. This is achieved by further refining the bounds until the first k_i upper and lower bounds are equal. While this can incur further computational effort, it is not equivalent to computing the complete $\mathbf{p}_i^{\langle t \rangle}$ vector as in Algorithm 2. The payments only need to be calculated for the units that are actually allocated on departure, the upper and lower bounds may be equal without needing to compute actual payments, and importantly, they are only required for the full set of agents and not recursively for all subsets of agents.

In practice, this algorithm with bounds significantly reduces the computational run-time of the mechanism (typically by 99% or more throughout the experiments conducted in Section 7), as it often avoids re-running the market with all possible subsets of agents. However, it is important to note that the worst-case run-time is still equivalent to Algorithm 2, i.e., $O(N \cdot 2^N)$, or exponential in the number of agents. In the best case, when no recursion is necessary, the run-time reduces to $O(N^2)$.

7. Experimental Evaluation

In this section, we quantify the performance of our mechanisms, as compared to a number of benchmarks, by applying them to a range of settings. While we investigated the theoretical performance bounds of our mechanisms in Section 5, the purpose of this section is to evaluate their performance in realistic settings.

Specifically, in Section 7.2, we consider a general setting that is easily reproducible and show how our mechanisms perform as we vary both supply and demand for a good. Then, in Section 7.3, we turn to the PHEV domain. For this, we first show how we can derive an agent's preferences based on the vehicle owner's driving behaviour. Then, we use real data collected during the first large-scale trial of pure electric vehicles (EVs) in the UK to show that the same trends continue to hold in a realistic application setting. Furthermore, we look at how the gradual introduction of fast-charging PHEVs would affect a neighbourhood with limited electricity supply, both in terms of social welfare (which translates to the overall fuel savings within the neighbourhood) and the financial savings of individuals. Throughout the experiments, we also consider a simpler greedy allocation mechanism without cancellation and quantify the potential benefits an agent would be able to achieve by misreporting in such a mechanism. This demonstrates whether there is actually scope for strategic misreporting in realistic settings and whether cancellation is needed in practice.

Before we consider the two specific settings, we briefly outline the common parameters and benchmarks used in all experiments.

7.1 Experimental Setup

To evaluate our mechanisms, we simulate different settings where a number of agents compete for a limited supply of a good that is allocated on an hourly basis over a 24-hour period. In order to test scenarios with varying supply and demand, we sample these agents randomly from fixed probability distributions and use a range of supply functions (these are outlined in more detail in Sections 7.2 and 7.3). In order to ensure statistical significance of our results, we re-sample the agents 1,000 times for each setting, and we plot 95% confidence intervals throughout this section.

In addition to the two mechanisms proposed in this paper, with immediate cancellation (`IM`) and on-departure cancellation (`OD`), we evaluate a number of benchmark mechanisms:

- `Fixed` is a fixed-price mechanism that allocates units to those agents that have a valuation of at least a given constant p, and the price they pay for each unit is p. When demand exceeds supply, each unit is allocated to an agent chosen uniformly at random from the set of all agents with a sufficiently high valuation. Here, an agent may receive multiple units, up to its maximum consumption rate, $\hat{r}_i$. This mechanism is DSIC and so it constitutes a direct comparison to our mechanisms. However, to optimise the performance of the fixed-price mechanism, p must be carefully chosen. Thus, for each given setting, we test all possible values (in steps of £0.01) and select the p that achieves the highest average efficiency (over 1,000 trials). Thus, when showing the results of `Fixed`, this constitutes an upper bound of what could be achieved with this mechanism.

- `Random` is a special case of `Fixed`, with $p = 0$. Thus, using this baseline benchmark, units are allocated randomly and agents do not pay anything.

- `Greedy` is a simple greedy allocation policy, as described in Section 4. Payments are calculated using $\mathbf{p}_i$ prices (as for `IM` and `OD`), but there are no cancellations. Thus, this mechanism is not truthful, but it constitutes an interesting comparison to our mechanisms, as it allows us to quantify the loss of efficiency that is caused by cancellations, as well as the potential benefits an agents has when misreporting in the absence of cancellations.

- `Heuristic` allocates units such that a weighted combination of an agent's valuation and urgency (proximity to its departure time) is maximised. Here, an $\alpha \in [0, 1]$ parameter denotes the importance of the urgency, such that $\alpha = 1$ corresponds to the well-known earliest-deadline-first heuristic in scheduling (Pinedo, 2008), while $\alpha = 0$ indicates that units are always allocated to the agent with the highest valuation. This is not a truthful mechanism and we do not impose payments here, as its primary purpose is as a benchmark for our approach. Again, we always select the best α by testing all values in steps of 0.01 for each setting.

- `Optimal` assumes complete knowledge of all future arrivals and supply, then allocates units to agents to maximise the overall allocative efficiency. Clearly, this mechanism is not possible because it assumes knowledge of the future and it is also not truthful (again we impose no payments), but it serves as an upper bound for the efficiency that could be achieved.

7.2 General Allocation Setting

First, we consider a general synthetic setting, in which we generate agents and the supply function from simple distributions. The main reason for examining such a scenario before turning to a more realistic setup is to generate results that are easily reproducible and that are not tied to a specific application domain. In the following, we outline the distributions from which we sample the supply and agents (Section 7.2.1), and then we discuss our results (Section 7.2.2).

7.2.1 Synthetic Setup

In this setting, we generate the supply function $S(t)$ by randomly drawing from the discrete uniform distribution on $\{1, 2, 3, \ldots, s\}$, where we vary s in our experiments to represent different amounts of a good that is being produced. For each agent i, we sample its arrival time a_i from the discrete uniform distribution on $\{0, 1, 2, \ldots, 23\}$ and its departure time from $\{a_i, a_i + 1, \ldots, 23\}$. We sample its maximum consumption rate r_i from $\{1, 2, 3, 4, 5\}$, and finally, we generate $\mathbf{v}_i$ by first selecting a number of required units uniformly at random from $\{1, 2, 3, \ldots, 20\}$. Then, the first valuation $v_{i,1}$ is sampled from an exponential distribution with rate 1, and the remaining valuations are drawn uniformly at random from the continuous interval $[0, v_{i,1}]$ (ordered appropriately to ensure non-increasing marginal valuations).

7.2.2 Synthetic Results

In Figure 3, we examine the allocative efficiency of the mechanisms as we increase the number of agents competing for a limited supply of electricity. The figure shows the allocative efficiency both of a setting with low supply (left), where $s = 1$, i.e., one unit is available per time step, and of a setting with high supply (right), where $s = 20$, i.e., up to 20 units are available per time step. We choose these two extreme settings to show the full spectrum of potential supply scenarios (and focus on supply settings based on real data in Section 7.3). Note that due to its run-time complexity, we plot `IM` only in the smaller setting with one unit of supply. When supply is very high ($s = 20$), each agent is typically allocated a large number of units, causing `IM` to require more frequent refinements of the bounds for the $\mathbf{p}_i^{\langle t \rangle}$ vectors and thus leading to a computational bottleneck. As the non-truthful `Heuristic` approach consistently achieves around 99% of the `Optimal` with full information when $s = 1$, we do not plot it for readability, and we also use it as an approximation of the `Optimal` when $s = 20$ (where `Optimal` also becomes computationally infeasible).

Several trends emerge in these results. First, when $s = 1$, the simple but not truthful `Greedy` approach performs very well (around 99% of the `Optimal`). Next, we note that both our truthful mechanisms, `IM` and `OD` also perform well, achieving around 95% and 96% of the

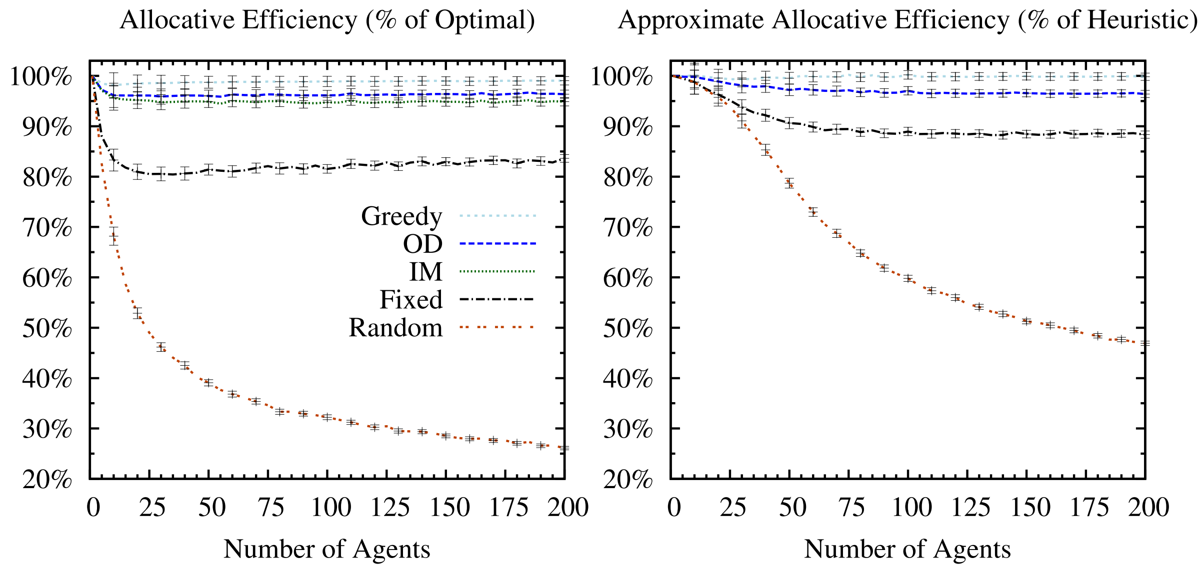

Figure 3: Allocative efficiency in synthetic setting (low supply, with $s = 1$, on the left, and high supply, with $s = 20$, on the right).

`Optimal`, respectively. This is slightly lower than `Greedy`, which indicates that about a 3–4% loss in efficiency is incurred due to cancellations. The difference in performance between `IM` and `OD` is as expected here, as `IM` can cancel more units (see Section 5); however, despite the unfavourable worst-case performance of `IM`, it is surprising that the difference is not significant in practice. Overall, the results are promising, indicating that our mechanisms work well in these settings, because the specific conditions that cause cancellations (i.e., when the valuations of allocated units effectively cross over with the marginal payments) do not occur frequently in practice.

The fixed price mechanism, `Fixed`, performs significantly worse than our proposed mechanisms, achieving only 81% – 83% of the `Optimal`. This is because, in order to remain truthful, the mechanism sets a single fixed price that does not respond dynamically to changes in supply and demand from time step to time step. This also explains why the mechanism performs worst when there is some, but not much, competition, e.g., around 30 agents. Here, the fixed price starts to rise, to ensure agents with higher valuations are allocated first, but there is still considerable variance in the valuations at each time step, sometimes leaving high-value agents unallocated, while other times units are not allocated at all. In contrast to this, both `IM` and `OD` always allocate all available units to the agents with the highest valuations. Furthermore, it should be noted that `Fixed`, unlike `IM` and `OD`, assumes a priori knowledge of the distributions from which agents are drawn. This may not always be available in practice, which can further decrease its performance.

Finally, the `Random` mechanism performs worst of all, which is not surprising, as it uses no information about the agents' valuations at all. However, its poor performance demon-

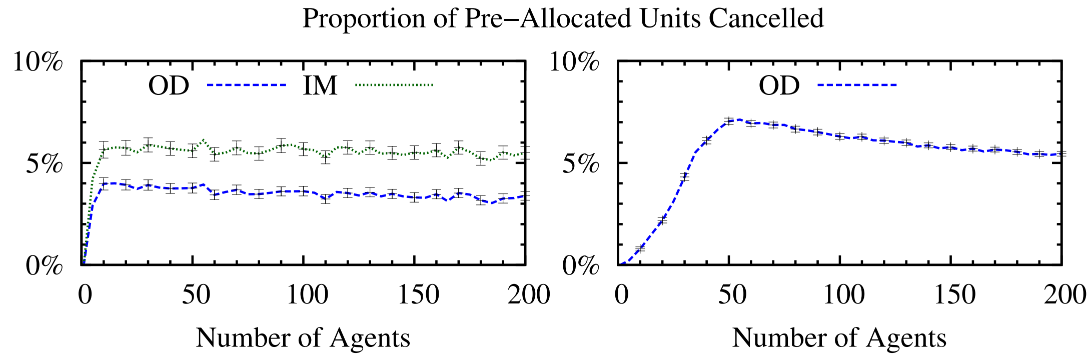

Figure 4: Proportion of initially allocated units cancelled ($s = 1$ on the left, and $s = 20$ on the right).

strates clearly the potential perils of using poorly designed non-truthful mechanisms, where strategically misreported valuations may have no relation to the agents' actual valuations.

In the setting with more abundant supply, when $s = 20$, the same broad trends are observed. The `OD` mechanism still achieves around 97% of the near-optimal `Heuristic` and slightly less than `Greedy`, while the `Fixed` and `Random` perform significantly worse. However, the gap in performance is smaller this time, as there is less competition and often there are sufficient units to satisfy most of the agents. In such settings, the benefit of always allocating to the agents with the highest valuations is generally lower. Note that while `IM` was infeasible in these settings due to its high run-time complexity, `OD` took, on average, less than 100ms to execute all 24 time steps in even the most complex settings with $s = 20$ and 200 agents.

Next, to further illustrate that actual cancellations by the `IM` and `OD` mechanisms are far from the worst-case bounds established in Section 5, Figure 4 shows the average proportion of pre-allocated units cancelled. This is generally low, ranging from 0% to 7%. As expected, the `OD` mechanism cancels fewer units than the `IM` mechanism. Furthermore, there is a general trend to cancel few units when competition is low (as there are usually sufficient units to satisfy all agents, leading to mostly 0 valuations in the $\mathbf{p}_i^{\langle t \rangle}$ vectors). Cancellations then peak for medium levels of competition, after which they start to drop again slightly. This peak can be explained by the large variations in the $\mathbf{p}_i^{\langle t \rangle}$ vectors in these settings and also because agents are generally allocated more units than in settings with more competition (leading similarly to a higher variation in the valuations of an agent's allocated units).

In Figure 5, we next consider the potential benefits of misreporting when no cancellations are used (as in the `Greedy` mechanism). We measure this by computing the utility of each agent had on-departure cancellations been used and compare this to the actual utility gained in the `Greedy` mechanism. This constitutes the best deviation a single agent could have achieved with perfect hindsight of the actual $\mathbf{p}_i$ prices. Here we plot the proportion of cases where an agent can achieve any gain from misreporting (light blue), the conditional

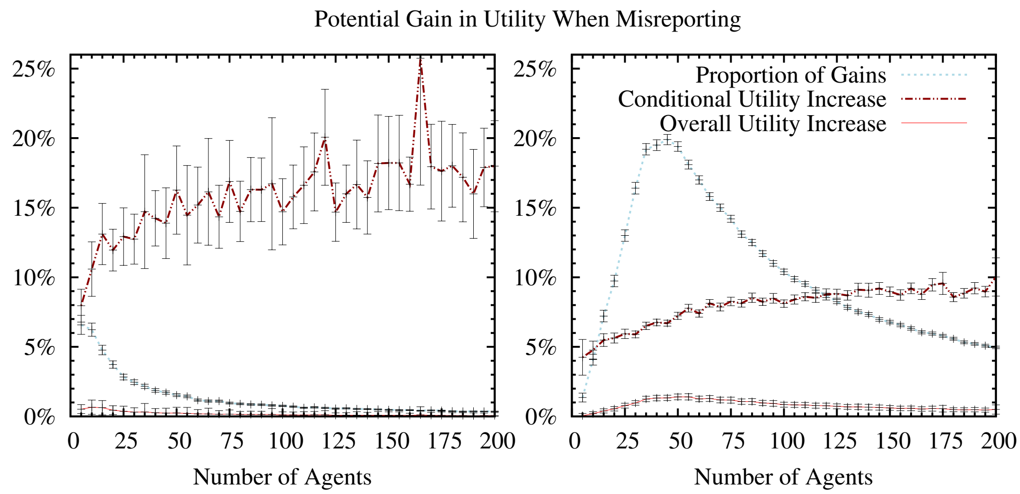

Figure 5: Potential gains when misreporting in the `Greedy` mechanism ($s = 1$ on the left, and $s = 20$ on the right).

proportional increase in utility when a gain from misreporting exists (dark red), and the product of the two, i.e., the overall average proportional increase in utility, including cases where there is no gain (light red).

The results indicate that without cancellations, there are often cases where an agent can benefit from misreporting — in up to 6–7% of all cases when $s = 1$ and up to 20% of cases when $s = 20$. This provides a clear motivation for using incentive compatible mechanisms in these settings. However, although individual gains can lead to an average increase in utility of up to 15–20%, when considering the overall average utility increase (including cases where agents do not benefit), this is only up to 1–2% and significantly less in many specific settings. This offers some promise for settings where cancellations are infeasible, for example because immediate cancellations are not computationally feasible in very large settings, or because on-departure cancellations cannot practically be implemented. For low expected gains of 1–2% or less, an agent may not wish to exert additional efforts to strategise. Furthermore, these gains represents an upper bound of what can be achieved with perfect foresight of the prices, which is likely to be unavailable in practice. We also note that the expected gain from misreporting fluctuates significantly, depending on the specific setting — similar to the cancellations, this fluctuation is caused by variations in the $\mathbf{p}_i$ vectors and also in the valuations of the allocated units.

To conclude this section, Figure 6 further explores the performance gap between our proposed mechanisms and the benchmarks as supply is increased. Here, we fix the number of agents at 50 and then vary the maximum number of available units of the good per time step, s, from 1 to 70. Due to the complexity of this setting, we again omit `Optimal` and `IM` from our analysis. It is clear here that the relative benefit the `OD` mechanism has over the other truthful benchmarks decreases as supply is increased. This is not surprising, as eventually all agents can be completely satisfied, even in a random allocation policy.

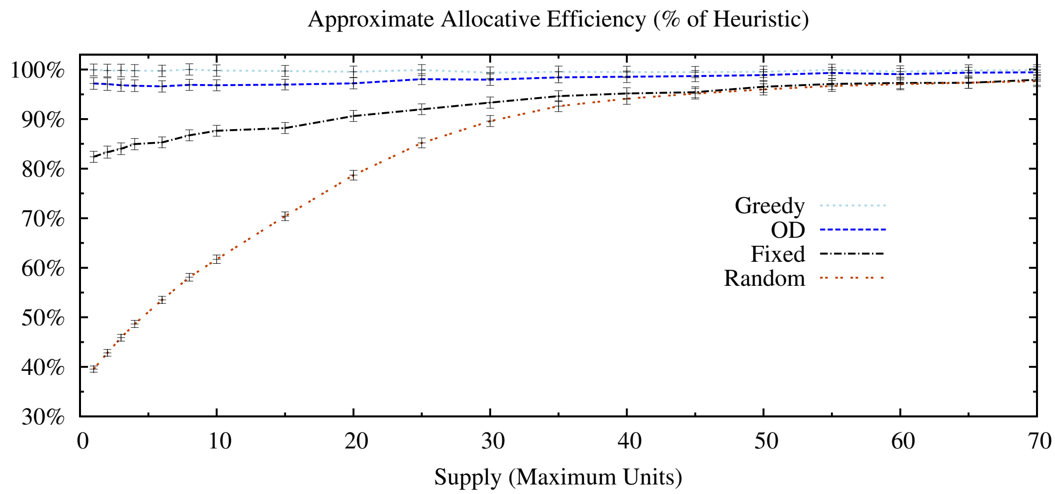

Figure 6: Approximate allocative efficiency as the maximum supply per time step (s) is increased.

However, there is still a significant benefit in using `OD` up to relatively high supply levels of around $s = 45$, and it is never outperformed by the truthful benchmarks, while consistently performing close to the near-optimal `Heuristic` (over 99% in some cases) and close to the `Greedy` mechanism without cancellations. In terms of potential gains for deviations, similar trends as discussed previously are observed and so we omit a detailed figure.

So far, we have concentrated on describing the general performance of our mechanisms in a synthetic, easily reproducible setting. In the following, we apply it to data from a real EV setting.

7.3 PHEV Setting

In this section, we use data from a real (pure) EV trial to simulate typical charging patterns.[8] Doing this allows us to verify that the trends discussed in the previous section continue to hold in a realistic setting. Furthermore, basing our experiments on actual PHEV characteristics enables us to quantify the actual utility of drivers in real terms (i.e., as a monetary gain or the fuel saved). We will also investigate whether the introduction of faster charging speeds will lead to any benefits in the settings we consider. Such an investigation is interesting, because fast chargers are already available for domestic settings and allow vehicles to charge at twice the normal rate (or faster).[9] However, their impact in our settings is unclear, as they are still constrained by the overall supply of electricity.

In the following, we first present a principled approach for deriving an agent's marginal valuation vector and show that this approach satisfies the non-increasing marginal valuation assumption (Section 7.3.1). We then describe the real-world data that we use for our

8. Note that, whereas the simulation is based on PHEVs, it uses real-world experimental data from pure EVs. However, we believe it is reasonable to assume that the charging behaviour would be similar.
9. See, for example, `http://www.pod-point.com/` or `http://www.charging-solutions.com/`.

experiments (Section 7.3.2), followed by an outline of how this is used to sample PHEVs (Section 7.3.3). Finally, we discuss our results (Section 7.3.4).

7.3.1 Deriving an Agent's Marginal Valuation Vector

An important part of our overall model is a method for computing the marginal valuation vector, $\mathbf{v}_i$, based on real data. To do so, we combine data about the sampled car's actual journey distances with a principled approach for calculating the expected economic benefit of charging for PHEVs. In detail, we first derive a probability density function, $p(m)$, from the data, which describes the probability of the distance travelled to be m miles (described in Sections 7.3.2 and 7.3.3). Given this distribution, the price of fuel (in £/litre), p_p, the internal combustion engine efficiency (in miles/litre), e_p, and the efficiency of the electric engine (in miles/kWh), e_e, we can then calculate the expected utility of a certain amount of charge (in kWh), c_e, as follows:

$$\mathbb{E}(u(c_e)) = \int_0^\infty \frac{p_p}{e_p} \cdot m \cdot p(m) \mathrm{d}m - \int_{c_e \cdot e_e}^\infty \frac{p_p}{e_p} \cdot (m - c_e \cdot e_e) \cdot p(m) \mathrm{d}m, \tag{9}$$

where the first term is the expected fuel cost without any charge, and the second term is the expected cost with a battery charge of c_e. Therefore, the utility function represents the *expected savings* in terms of real money for a given battery charge (without taking the cost of the charge into account). Given this, and a unit size (in kWh), s_e, it is straight-forward to calculate the marginal valuation of the k^{th} unit as follows:

$$v_k = \mathbb{E}(u(k \cdot s_e)) - \mathbb{E}(u((k-1) \cdot s_e)) \tag{10}$$

Recall that, in our model, we assume that valuations are marginally non-increasing. We now show that, using the above approach, this assumption is automatically satisfied. To do so, we need to show that Equation 9 is non-decreasing (i.e., the first derivative is positive) and concave (i.e., the second derivative is negative). The first derivative is given by:

$$\frac{\mathrm{d}\mathbb{E}(u(c_e))}{\mathrm{d}c_e} = \frac{p_p}{e_p} \cdot e_e \cdot \int_{c_e \cdot e_e}^\infty p(m) \tag{11}$$

The second derivative is given by:

$$\frac{\mathrm{d}^2\mathbb{E}(u(c_e))}{\mathrm{d}{c_e}^2} = -\frac{p_p}{e_p} \cdot {e_e}^2 \cdot p(c_e \cdot e_e) \tag{12}$$

Clearly, both conditions are satisfied, which means that the valuations are always positive and marginally non-increasing. In what follows, we describe how we derive the experimental settings, such as the supply function, the arrival and departure of agents, and the travel distance probability distributions from a real-world dataset. We also provide examples of the marginal valuation vectors using this data.

7.3.2 The CABLED Dataset

We base our experiments on data gathered by the CABLED (Coventry And Birmingham Low Emissions Demonstration) project,[10] which is the first large-scale endeavour in the

10. See `http://cabled.org.uk/`.

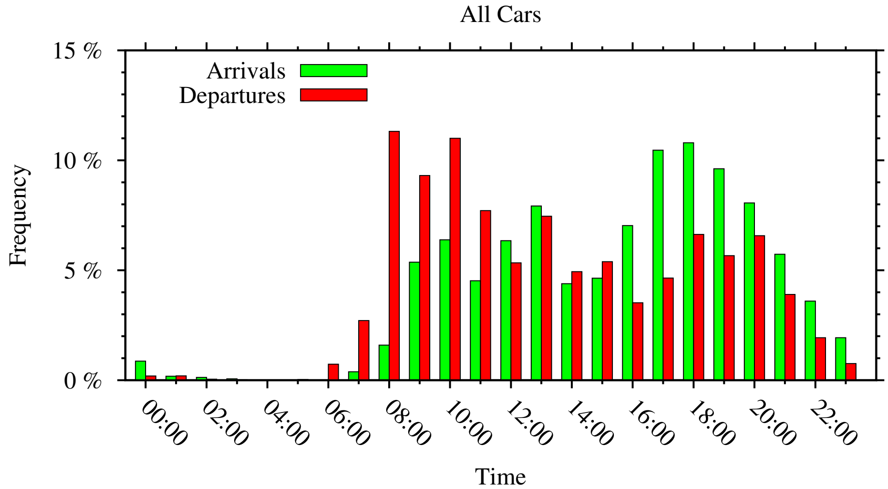

Figure 7: Distributions of arrival and departure times for all 56 EVs in the CABLED dataset (assigning equal weight to each EV).

UK to record and study the driving and charging behaviours of EV owners. As part of this project, 110 EVs were loaned to the public and equipped with GPS and data loggers to record comprehensive usage information, such as trip durations and distances, home charging patterns and energy consumption.

From this data, we focus on the period from March to June 2011, for which we were provided with information from 63 distinct vehicles, with a total of 13,273 journeys. For each journey, this includes the times when the ignition was turned on and off again, the total mileage (as derived from GPS readings that were taken every 60 seconds), as well as labels for the starting and end location, when this was available (such as "home" or "work").

The vehicles in the CABLED trial were charged at various locations — mostly at home, but also at work. For the purpose of our experiments, we will assume that all charging of the vehicles takes place at a single location. This is because our work focuses on coordinating the charging of EVs within a specific neighbourhood and considering the effect of multiple markets for electricity is beyond the scope of this work. When available, we choose this charging location to be the one labeled "home" in the data.[11] Given this, and since we are only interested in the arrival and departure times at this charging location, as well as the consumption patterns between visits to the charging location, we aggregate all intermediate journeys between the departure from a charging location until the next return to this location into a single journey.

Aggregating the data in this way and discarding vehicles without a clear charging location results in 4,302 distinct journeys for 56 different EVs, covering a total distance travelled of close to 72,500 miles. The overall distribution of all recorded arrival and departure times

11. Some vehicles in the dataset lack this, as they were used as shared fleet vehicles for an organisation — in these cases, we use an appropriate alternative location label, where most of the charging took place, or discard the vehicle when no suitable label can be identified.

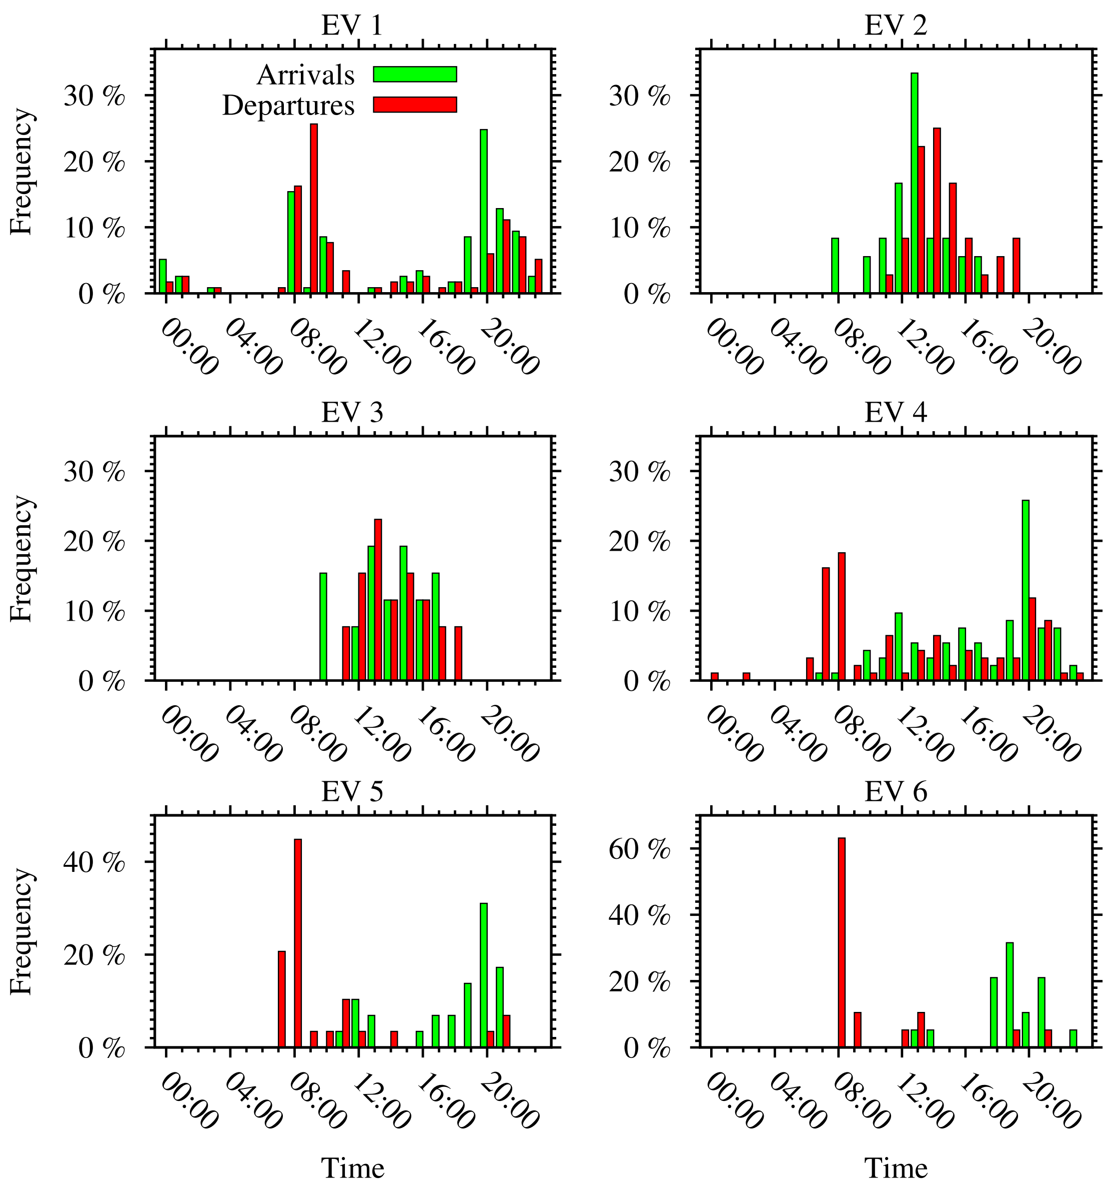

Figure 8: Example distributions of arrival and departure times for six EVs.

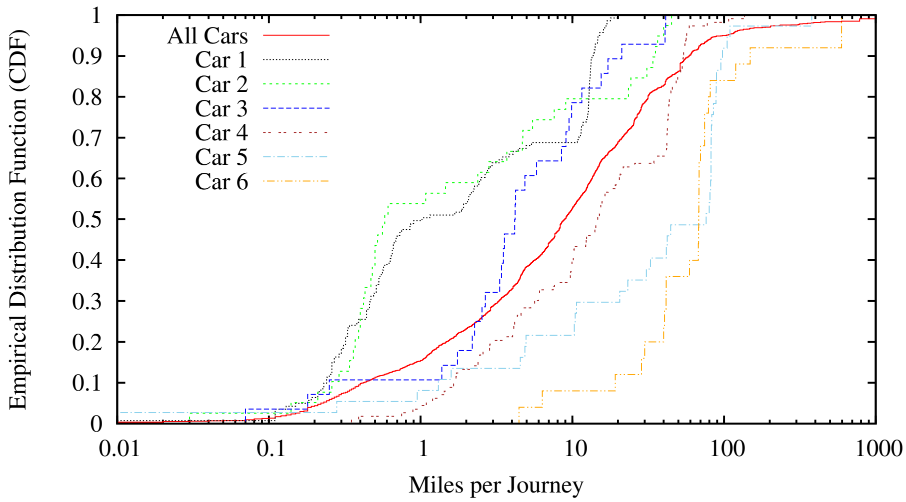

Figure 9: Empirical distribution functions of distances travelled between successive visits to home charging locations. Six example EVs and the combined summary of all 56 EVs, with equal weights assigned to all cars (in red), are shown.

at the charging location is shown in Figure 7.[12] Here, almost all arrivals happen between 9am and 11pm, with a clear peak in the late afternoon and evening (4pm – 9pm). Departures take place throughout the day with a peak around 8am – 10am. To show how individual driving patterns vary between the recorded cars, Figure 8 details the arrival and departure time distributions for six individual EVs from the data set. Most of these reflect the general trends shown in the previous figure, with arrivals generally occurring in the evening and most departures in the morning. However, EVs 2 and 3 deviate from this pattern. This is because those vehicles are shared fleet vehicles, which are collected and returned to their main charging location throughout the day.

Figure 9 shows the distribution of the distances the vehicles travelled between visits to their home charging location (red line), as well as the corresponding functions of the six EVs from Figure 8 (interrupted lines).[13] Overall, the average distance travelled is about 41 miles, while the median is around 9 miles (assigning uniform probabilities to each car type). The six sample EVs here show significantly different typical travel distances, ranging from an average of 4.19 miles (EV 1) to over 100 miles (EV 6).

In the following, we discuss how we use the data from the CABLED project to instantiate our EV charging simulations.

12. As we will allocate electricity in hourly units, all arrival and departure times are here rounded up to the next full hour.
13. Note that the distance of some journeys here exceeds the range of typical electric cars. This is because they were charged at alternative locations during the CABLED trial, which we ignore in our experiments. Since we focus on PHEVs in this work, in practice, the shortfall here would be made up by the combustion engine.

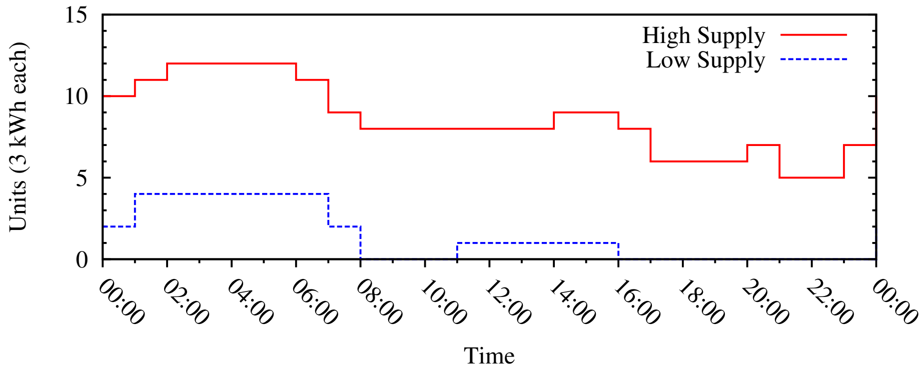

Figure 10: Number of 3 kWh units available for PHEV charging in two scenarios with varying supply (*high* and *low*), based on a neighbourhood of 30 households.

7.3.3 Generating Experiments

For each experimental run, we simulate a small neighbourhood of 30 households with a variable number of PHEVs over a 24-hour period, starting at 8:00 am in the morning until 8:00 am the following day. We assume that electricity is allocated in hourly time steps, where each unit corresponds to 3 kWh (which is the approximate energy obtained through a standard 13 A BS 1363 household socket in the UK when charging for an hour).

To obtain the supply function $S(t)$, we first compute the overall average electricity consumption throughout the neighbourhood (without PHEVs), based on the average consumption of a single UK household during a weekday in high summer.[14] Then, we assume that overall electricity supply is limited by the capacity of the local transformer, such that the electricity available for PHEV charging, $S(t)$, is the difference between this capacity (possibly including some additional safety margin) and the current overall consumption. In more detail, we consider two possible scenarios: (1) a *high supply* scenario, where the capacity limit is set such that it covers 150% of the peak consumption (at about 10:00 pm), resulting in 615 kWh available for PHEV charging; and (2) a *low supply* scenario, where the capacity limit is 80% of peak consumption, resulting in 99 kWh available for charging.[15] The corresponding units available for PHEV charging in these two scenarios are shown in Figure 10.

Furthermore, we use the specific empirical distribution of journey distances corresponding to that car's type as that car's travel distance distribution, $p(m)$ (for example, if the sampled car is based on car 3, we use the dashed dark blue distribution function in Figure 9). We then use Equations 9 and 10 to derive the marginal valuations. In this case, as the empirical distribution function is discrete, the integrals in Equation 9 are replaced by sums over the data points. Furthermore, we will initially assume that r_i is drawn uniformly

14. We used the data available at `http://data.ukedc.rl.ac.uk/browse/edc/Electricity/LoadProfile/data`.

15. Local transformers are often undersized in this way since prior to PHEV use, they could cool during overnight periods of low demand.

at random from the discrete set $r_i \in \{1, 2, 3, 4\}$, that is, cars can charge between one to four units of electricity per hour (corresponding to 3 – 12 kWh).

To generate agents with a variety of marginal valuations, we note that e_e and e_p depend on the specific make and type of the PHEV. To simulate this, we draw e_e uniformly at random from 2 – 4 miles/kWh and e_p is drawn from 9 – 18 miles/litre. Next, we draw the capacity of a car battery from 15 – 25 kWh. These are realistic values modelled on the Chevrolet Volt, the first mass-produced PHEV. However, we include some variance to account for other vehicle types. Throughout the experiments, we hold the price of petrol constant at $p_p = \pounds 1.30$ per litre.

Table 4 shows example valuations corresponding to the same six cars considered previously (fixing $e_p = 13.5$ miles/litre, $e_e = 3$ miles/kWh and the battery capacity at 20 kWh). These highlight how longer expected journeys generally translate to higher marginal valuations, but also how variable the valuations can be for an individual agent. As an example, car 4 values the first 3 kWh of electricity at £0.67, but the seventh unit is only worth £0.038, as it is far less likely to be used.

Car	$v_{i,1}$	$v_{i,2}$	$v_{i,3}$	$v_{i,4}$	$v_{i,5}$	$v_{i,6}$	$v_{i,7}$
1	0.340	0.136	0.001				
2	0.304	0.178	0.162	0.114	0.033		
3	0.481	0.157	0.073	0.062	0.035		
4	0.670	0.453	0.333	0.312	0.263	0.134	0.038
5	0.727	0.620	0.582	0.540	0.498	0.445	0.445
6	0.839	0.797	0.767	0.711	0.630	0.555	0.540
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

Table 4: Example marginal valuations (in £).

To set up an experimental run, we then randomly generate a set of N PHEV agents, where we vary N from 1 to 60 to simulate different levels of demand.[16] For each agent i, we first choose one of the 56 available cars uniformly at random from the CABLED dataset, which we base that agent's type on. Then, we randomly select one of the car's recorded journeys and use the time of day of the car's arrival at the charging location as a_i (at or after 8:00 am in the time window we consider). To ensure that the correlation between arrival times and subsequent departure times in the dataset are preserved, we use the departure time of the journey immediately following the sampled journey as d_i (or 10 hours after arrival, whichever is sooner).

7.3.4 PHEV Results

First, we are interested in general trends of the mechanisms and whether these are similar to the trends discussed in Section 7.2.2. To this end, Figure 11 shows the allocative efficiency for all mechanisms in a setting with low supply (where it is feasible to run both the `Optimal`

16. Note that this is a realistic number of PHEVs within a neighbourhood served by a single distribution transformer (Huang & Infield, 2010).

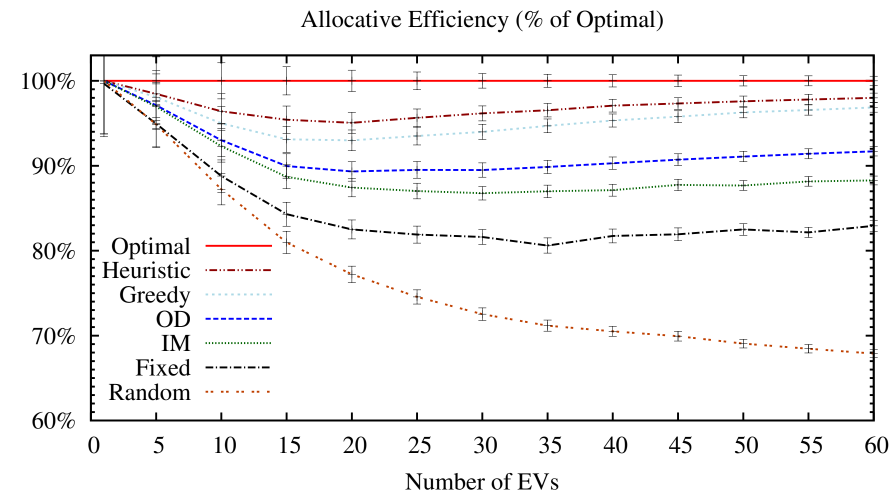

Figure 11: Allocative efficiency in a small neighbourhood of 30 households.

and the `IM` mechanisms). This demonstrates the same broad trends as in our previous synthetic setup — the `OD` and `IM` mechanisms clearly dominate the other truthful mechanism (with `IM` achieving slightly worse results due to its higher cancellation rates). This, again, is due to the ability of the mechanisms to always allocate to the agents with the highest valuations.

However, although they still consistently achieve around 90% of the `Optimal`, the relative performance of the `OD` and `IM` mechanisms here is slightly lower. This drop in performance, also witnessed by the `Heuristic` and `Greedy` mechanism, is due to the more constrained real-world settings, where electricity is only available in abundance at certain times (i.e., during the night), and where some agents are significantly less patient than others. In such settings, it can often pay off to delay more patient agents, even if they have higher valuations, in favour of less patient ones. Furthermore, because valuations are directly related to the fuel costs saved by a unit of electricity, there is less variance in the real-world valuations, causing the gap between the `OD` and `IM` mechanisms and the other truthful benchmarks to narrow slightly. Turning to the potential gains of misreporting in this setting, Figure 12 again confirms the same patterns observed previously. The magnitude of the gains are slightly higher here due to the different setting (reaching up to 3% in terms of overall gains and up to 30% in the conditional case).

A key advantage of applying the mechanisms to real-world data is that it allows us to determine the actual fuel savings agents could achieve in these settings. Thus, Figure 13 shows the average fuel savings of each agent under the various mechanisms, or, in other words, the average amount each agent would have spent on fuel, had they not been allocated any electricity. Initially, this is high (around £1.15), as there is little competition, but starts dropping as more PHEV owners compete for the same amount of electricity. Of key interest here is the horizontal separation between the different mechanisms. For a given fuel saving per agent, our mechanism can sustain a significantly larger number of agents than the other truthful mechanisms. For example, to save at least £0.40 per agent, `Random` can support

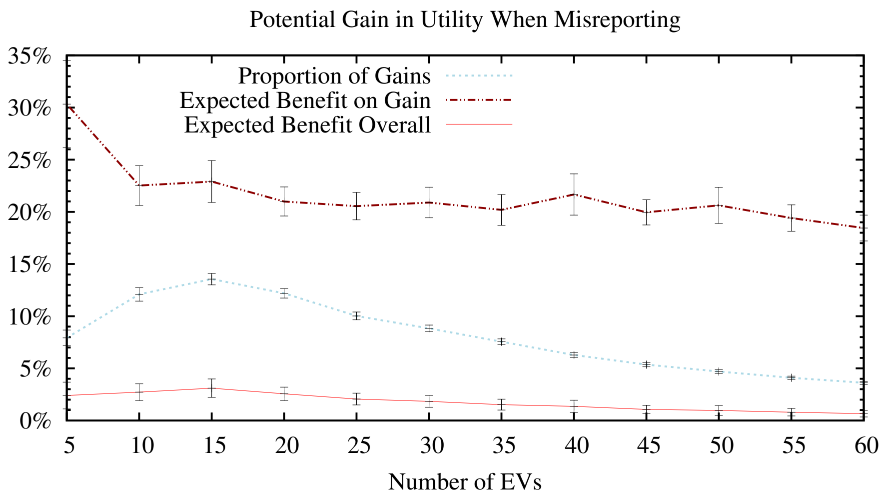

Figure 12: Potential gains when misreporting in the `Greedy` mechanism in a small neighbourhood of 30 households.

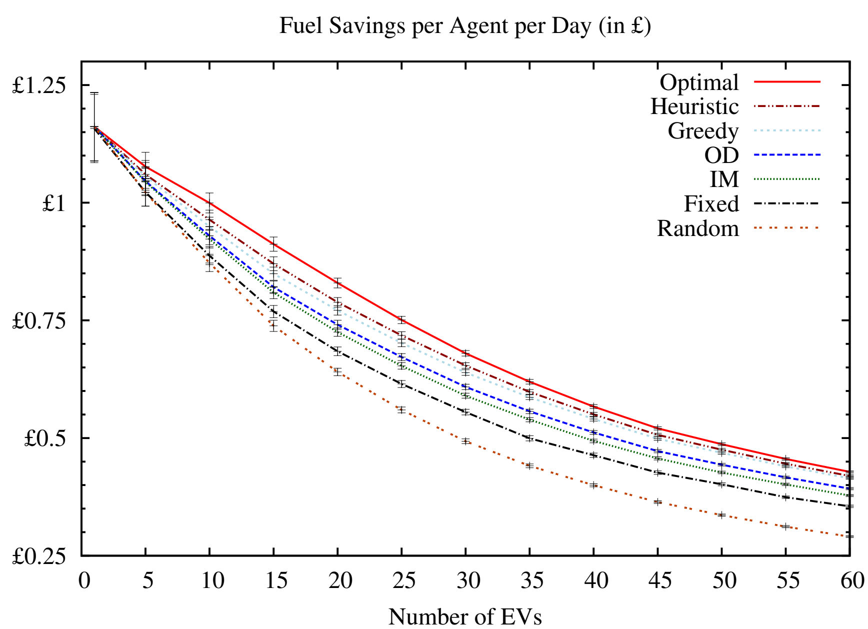

Figure 13: Saving per agent per day in a small neighbourhood of 30 households.

up to 40 PHEV owners, while `IM` and `OD` achieve the same threshold for around 60 PHEV owners (an approximately 50% improvement).

Finally, we consider in more detail how the presence of fast-charging vehicles affects the overall neighbourhood, in terms of overall fuel savings, the occurrence of cancellations and the utilities of individual agents. To this end, we now fix the number PHEVs at 60

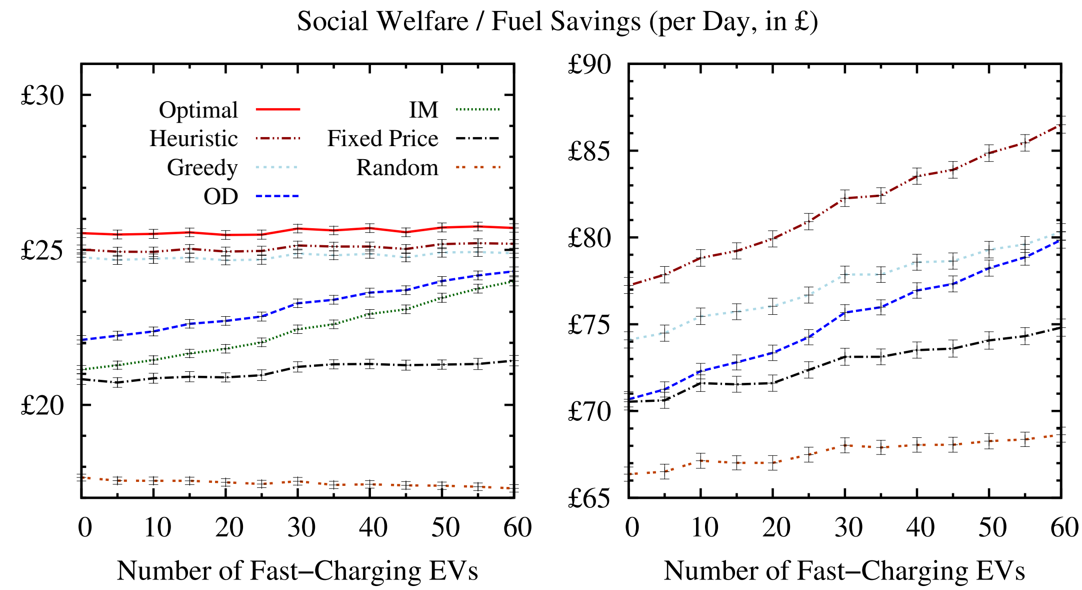

Figure 14: Social welfare when introducing fast-charging cars into a neighbourhood with low supply (left) and high supply (right).

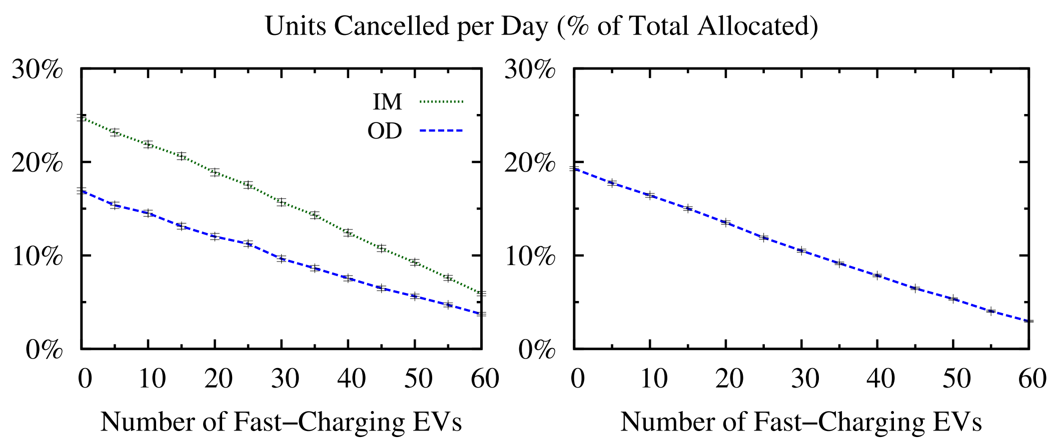

Figure 15: Cancellations when introducing fast-charging cars into a neighbourhood with low supply (left) and high supply (right).

and consider both the low and high demand settings shown in Figure 10. Due to their computational cost, we again only test the `Optimal` and `IM` mechanisms in the setting with low supply. To investigate the impact of fast-charging, we assume there are two agent types — the first, *normal*, can charge a single unit of 3 kWh per time step, while the second, *fast*, are equipped with fast chargers that can charge up to four such units per time step.

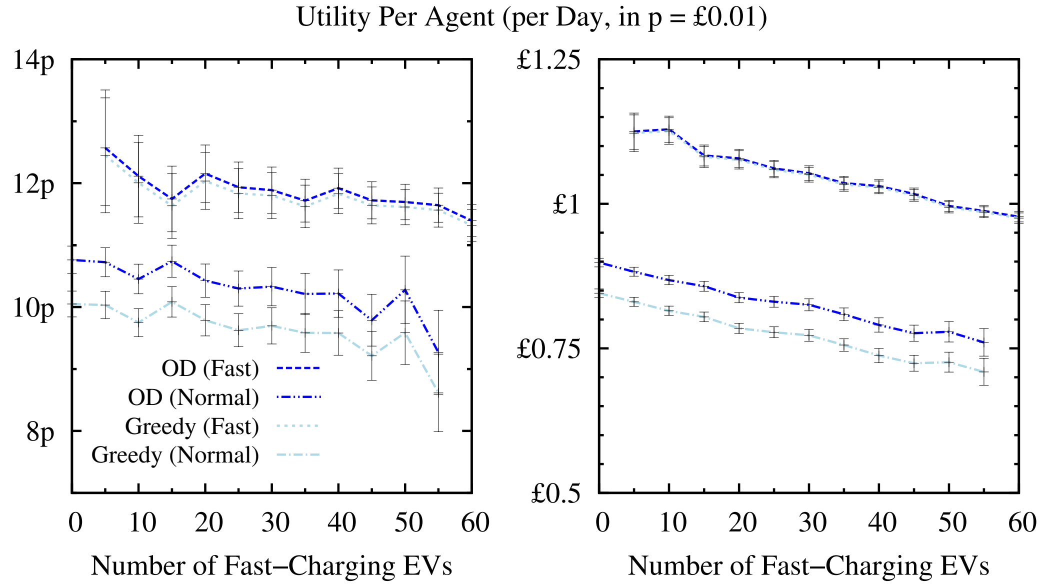

Figure 16: Individual agent utility when introducing fast-charging cars into a neighbourhood with low supply (left) and high supply (right).

Throughout the experiments, we vary the number of fast-charging PHEVs (out of the total 60). Figure 14 first shows the resulting social welfare (i.e., overall fuel savings) for both supply scenarios with low supply (left) and high supply (right). First, we note that the trends for the two scenarios are different – when supply is low, the introduction of more fast-charging vehicles has little effect on overall social welfare for most mechanisms, while when supply is high, most mechanisms display increased savings. This happens because the first scenario is highly constrained, with the low supply resulting in few occasions where an agent could charge more than a single unit per time step. In contrast, when supply is high, agents are often allocated multiple units, thus enabling impatient agents in particular to achieve higher overall fuel savings.

In addition to this, it is interesting to note that our proposed mechanisms `OD` and `IM` benefit in *both* settings (achieving additional fuel savings of almost two litres per day in the low supply setting, and up to seven litres in the high supply setting). The reason for this becomes evident when considering the proportion of units cancelled as more fast-charging PHEVs are introduced — for both mechanisms and in both settings, the number of units cancelled are consistently reduced by around 70–80% as all cars are replaced by fast-charging PHEVs (shown in Figure 15). This occurs mainly because there are more active marginal valuations at each time step to populate the $\mathbf{p}_i^{\langle t \rangle}$ vectors, thus reducing the number of cancellations. This also causes the gap between our mechanisms and the `Greedy` mechanism to shrink, as fewer cancellations take place.

With respect to the utility of individual agents (including payments to the mechanism), Figure 16 shows that agents *in both settings* always have an incentive to switch to fast-

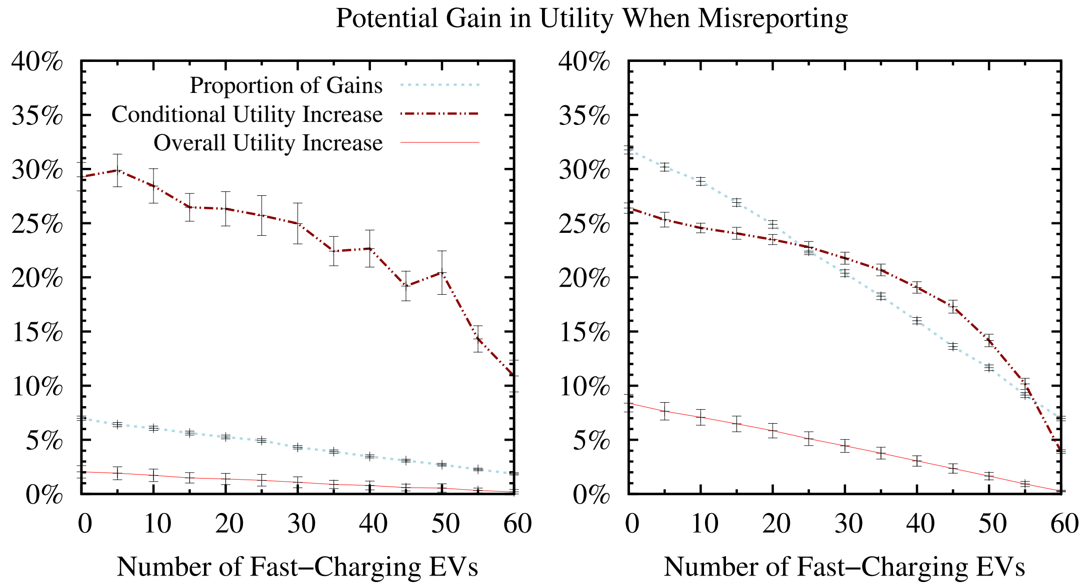

Figure 17: Potential gains when misreporting in the `Greedy` mechanism in a neighbourhood with low supply (left) and high supply (right).

charging PHEVs (e.g., by purchasing a domestic fast charger), and this applies for both the `OD` mechanism and the `Greedy` mechanism. With low supply, the expected daily saving when switching to a fast-charging PHEV is approximately £0.02–0.03, while with high supply, this is around £0.20–0.25. In both cases, this benefit is the result of increasing available supply per time step, as well as increasing the size of the price vector. Furthermore, even if the entire population were to switch from slow charging PHEVs to fast-charging PHEVs, individuals would, on average, achieve a higher utility. Note that the differences between the utility in the `OD` and `Greedy` mechanisms are significantly smaller for fast-charging vehicles, indicating that fast-charging agents can expect lower gains from misreporting when there are no cancellations.

Figure 17 further investigates the individual gains from misreporting when there are no cancellations. This shows an interesting trend — while initial gains are high (reaching over 8% in one setting), they decrease significantly as more fast-charging PHEVs are introduced (to around 0.2% in the same setting). This is clearly due to the significant reduction in cancellations that are witnessed in those settings. Furthermore, we note, by comparing `OD` and `Greedy` in Figure 16, that the agents who can gain from misreporting tend to be only the slow-charging ones. Overall, this is a promising result for settings where cancellations are not feasible — by increasing the consumption rate of PHEVs (within realistic parameters that are achievable by current technological trends), the scope and potential benefits from strategising in a simple greedy mechanism can be reduced significantly. In this particular example, when supply is high, a fast-charging PHEV can expect to gain less than £1 over the course of an entire year by strategising optimally.

8. Conclusions

The contributions in this paper are both theoretical and practical. On the theoretical side, we propose a novel online, model-free mechanism for perishable goods where agents have multi-unit demand and non-increasing marginal valuations. We show that, in order to ensure dominant strategy incentive compatibility in such a setting, our mechanism occasionally requires units to remain unallocated (we say their pre-allocation is cancelled), even if there is demand for these units. We define two ways in which cancellation can be performed: *immediate*, i.e., before the actual allocation, or *on departure* of the agent from the market. We study the properties of these two variants, both in terms of their incentives and allocative efficiency. Furthermore, we present algorithms for computing the payments and allocations of both mechanisms, and analyse their computational tractability.

The on-departure cancellation mechanism has better computational tractability, and has the same worst-case competitive bound, in terms of allocation efficiency, as the single-unit demand case. However, this mechanism requires any cancellations to be done on departure of an agent from the market which is not always feasible. A naïve approach to computing payments in the mechanism with immediate cancellations requires time that is exponential in the number of agents. To address this, we proposed a branch-and-bound algorithm that allows payments to be computed in the immediate cancellation policy for many realistically sized settings. Another potential problem with the immediate cancellation policy is that there is no worst-case bound in terms of the efficiency of the allocation.

On the practical side, we show how our mechanism can be applied within the smart grid to solve the important problem of integrating an increasing number of high-consumption PHEVs into the electricity grid. In addition to a synthetic setting, we empirically evaluate our mechanism using real-world data from a large-scale trial of electric vehicles in the UK. We show that the proposed mechanism is highly robust, scalable (in particular, the on-departure variant) and achieves better allocative efficiency than any fixed-price benchmark, while only being slightly less efficient than an established cooperative scheduling heuristic. Specifically, we demonstrate that our mechanism can sustain up to 50% more vehicles at the same fuel cost than can be achieved using a simple randomised mechanism. Both variants also consistently achieve an efficiency of around 90%, compared to a hypothetical optimal offline solution. Given the theoretical results regarding their bounds, this is a surprising result, suggesting that the specific conditions that cause cancellations do not often occur in practice and that our allocation policies perform well in realistic settings. Furthermore, we consider the introduction of fast chargers within a neighbourhood, and we show that this leads to a significant increase in overall fuel savings and that it further reduces the occurrence of cancellations. Finally, since on-departure cancellation requires discharging the battery of the PHEV, we consider the the potential gains from misreporting if the units are not cancelled (and assuming full knowledge of the types of other agents). From the settings we considered, the average potential gain is between 1% and 8% overall, but could go up to 30% on average when only considering those cases where misreporting is beneficial, and could be even higher for individual cases. The gain becomes smaller as the number of competing agents increases and, interestingly, fast-charging PHEVs have particularly low incentives for misreporting.

Taken together, our mechanisms represent a versatile range of tools, some of which may be more suitable for specific scenarios than others. For example, in medium-sized settings where allocations cannot be cancelled on departure, the `IM` mechanism may be the most suitable (e.g., in the low-supply PHEV settings outlined in Section 7.3). In other settings where on-departure cancellation is feasible, the `OD` mechanism leads to a higher average and worst-case efficiency, and is also more scalable. Here, it is also important to emphasize that on-departure cancellation only occurs when this is in the user's best interest— thus, it is entirely possible to achieve through an optional action. Finally, as we show in the results, even when both `IM` and `OD` are infeasible, a mechanism without cancellations may still be viable in some settings, and it may even be possible to significantly reduce the scope for manipulation by adjusting some of the system parameters; e.g., by introducing fast chargers in the PHEV setting or by increasing supply.

There are several directions for extending this work. In related work, Stein et al. (2012) discuss an alternative model, which uses probabilistic information about future arrivals and is designed to elicit truthful reporting from pure EVs, rather than PHEVs. That model, however, requires knowledge about future supply and assumes single-minded bidders; i.e., the preferences are single-dimensional and so it is not possible to specify different values for different amounts of charge received. In future work, we intend to explore mechanisms that combine the benefits of both approaches.

In addition, we intend to test our mechanism using a real-world trial. As we have seen (see Section 7.3), we can design an agent which elicits information regarding the intended use (a probability distribution of the driving distance), and combines this with other information such as the price of petrol and the efficiency of the vehicle, to derive the owner's marginal valuation vector. Such an agent can also participate in the mechanism on the owner's behalf, avoiding the need for the owner to understand the details of the mechanism. This is in the spirit of the work on *hidden market design* (Seuken, Parkes, Horvitz, Jain, Czerwinski, & Tan, 2012), where the aim is to design the user interface such that the user's cognitive load is reduced by hiding details of the underlying market. In the trial, we intend to approach participants owning regular (non EV) cars, and install GPS trackers in these cars. Participants will then be asked to predict their driving requirements, and their agent will use this information, as well as learned historic driving patterns, to derive the user's utility function and participate in the mechanism on their behalf. Although the trial will be with regular cars, the users will be able to see how much they would have hypothetically saved by providing accurate (and truthful) information about their intended use.

Acknowledgments

This work was supported by the iDEaS (*www.ideasproject.info*) and ORCHID (*www.orchid.ac.uk*) projects at the University of Southampton. David Parkes was supported in part through the Harvard SEAS TomKat fund.

References

Babaioff, M., Blumrosen, L., & Roth, A. (2010). Auctions with online supply. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC'10)*, pp. 13–22.

Bent, R., & Van Hentenryck, P. (2004). The value of consensus in online stochastic scheduling. In *Proceedings of the 14th International Conference on Automated Planning and Scheduling (ICAPS'04)*, pp. 219–226.

Bikhchandani, S., Chatterji, S., Lavi, R., Mu'alem, A., Nisan, N., & Sen, A. (2006). Weak monotonicity characterizes deterministic dominant-strategy implementation. *Econometrica*, *74*(4), 1109–1132.

Blum, M., Floyd, R. W., Pratt, V., Rivest, R. L., & Tarjan, R. E. (1973). Time bounds for selection. *Journal of Computer and System Sciences*, *7*(4), 448 – 461.

Clement, K., Haesen, E., & Driesen, J. (2009). Coordinated charging of multiple plug-in hybrid electric vehicles in residential distribution grids. In *Proceedings of the IEEE/PES Power Systems Conference and Exposition (PSCE'09)*, pp. 1–7.

Constantin, F., Feldman, J., Muthukrishnan, S., & Pal, M. (2009). An online mechanism for ad slot reservations with cancellations. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA'09)*, pp. 1265–1274.

Constantin, F., & Parkes, D. C. (2009). Self-correcting sampling-based dynamic multi-unit auctions. In *Proceedings of 10th ACM Conference on Electronic Commerce (EC'09)*, pp. 89–98.

Eberle, U., & von Helmolt, R. (2010). Sustainable transportation based on electric vehicle concepts: a brief overview. *Energy & Environmental Science*, *3*, 689–699.

Engel, Y., & Wellman, M. P. (2010). Multiattribute auctions based on generalized additive independence. *Journal of Artificial Intelligence Research (JAIR)*, *37*, 479–525.

Fairley, P. (2010). Speed bumps ahead for electric-vehicle charging. *IEEE Spectrum*, *47*(1), 13–14.

Gerding, E., Stein, S., Robu, V., Zhao, D., & Jennings, N. R. (2013). Two-sided online markets for electric vehicle charging. In *Proceedings of 12th International Confernece on Autonomous Agents and Multiagent Systems (AAMAS'13)*, pp. 989–996.

Gerding, E., Robu, V., Stein, S., Parkes, D., Rogers, A., & Jennings, N. (2011). Online mechanism design for electric vehicle charging. In *Proceedings of the 10th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'11)*, pp. 811–818.

Gershkov, A., & Moldovanu, B. (2010). Efficient sequential assignment with incomplete information. *Games and Economic Behavior*, *68*(1), 144–154.

Hajiaghayi, M., Kleinberg, R., Mahdian, M., & Parkes, D. C. (2005). Online auctions with re-usable goods. In *Proceedings of 6th ACM Conference on Electronic Commerce (EC'05)*, pp. 165–174.

Huang, S., & Infield, D. (2010). The impact of domestic plug-in hybrid electric vehicles on power distribution system loads. In *Proceedings of the International Conference on Power System Technology (POWERCON 2010)*, pp. 1–7.

Kamboj, S., Kempton, W., & Decker, K. S. (2011). Deploying power grid-integrated electric vehicles as a multi-agent system. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'11)*, pp. 13–20.

Lavi, R., & Nisan, N. (2004). Competitive analysis of incentive compatible on-line auctions. *Theoretical Computer Science*, *310*, 159–180.

Nisan, N., Roughgarden, T., Tardos, E., & Vazirani, V. (2007). *Algorithmic Game Theory*. Cambridge University Press.

Parkes, D. C. (2007). Online mechanisms. In Nisan, N., Roughgarden, T., Tardos, E., & Vazirani, V. (Eds.), *Algorithmic Game Theory*, pp. 411–439.

Parkes, D. C., & Duong, Q. (2007). An ironing-based approach to adaptive online mechanism design in single-valued domains. In *Proceedings of the 22nd National Conference on Artificial Intelligence (AAAI'07)*, pp. 94–101.

Parkes, D. C., & Singh, S. (2003). An MDP-based approach to online mechanism design. In *Proceedings of the 17th Conference on Neural Information Processing Systems (NIPS'03)*, pp. 791–798.

Pinedo, M. (2008). *Scheduling: Theory, Algorithms, and Systems* (3rd edition). Springer.

Porter, R. (2004). Mechanism design for online real-time scheduling. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC'04)*, pp. 61–70.

Robu, V., Stein, S., Gerding, E., Parkes, D., Rogers, A., & Jennings, N. (2011). An online mechanism for multi-speed electric vehicle charging. In *Proceedings of the 2nd International Conference on Auctions, Market Mechanisms and their Applications (AMMA'11)*, pp. 100–112.

Robu, V., Kota, R., Chalkiadakis, G., Rogers, A., & Jennings, N. R. (2012). Cooperative virtual power plant formation using scoring rules. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI'12)*.

Robu, V., Noot, H., La Poutré, J. A., & van Schijndel, W. (2011). A multi-agent platform for auction-based allocation of loads in transportation logistics. *Expert Systems with Applications*, *38*(4), 3483–3491.

Sandholm, T. (2002). Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, *135*(1-2), 1–54.

Seuken, S., Parkes, D. C., Horvitz, E., Jain, K., Czerwinski, M., & Tan, D. (2012). Market user interface design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pp. 898–915. ACM.

Stein, S., Gerding, E., Robu, V., & Jennings, N. R. (2012). A model-based online mechanism with pre-commitment and its application to electric vehicle charging. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'12)*, pp. 669–676.

Stein, S., Gerding, E., Rogers, A., Larson, K., & Jennings, N. R. (2011). Algorithms and mechanisms for procuring services with uncertain durations using redundancy. *Artificial Intelligence*, *175*, 2021–2060.

Sundstrom, O., & Binding, C. (2012). Flexible charging optimization for electric vehicles considering distribution grid constraints. *IEEE Transactions on the Smart Grid*, *3*(1), 26–37.

Vasirani, M., & Ossowski, S. (2011). A computational monetary market for plug-in electric vehicle charging. In *Proceedings of the 2nd International Conference on Auctions, Market Mechanisms and their Applications (AMMA'11)*, pp. 88–99.

Vytelingum, P., Voice, T., Ramchurn, S. D., Rogers, A., & Jennings, N. R. (2011). Theoretical and practical foundations of large-scale agent-based micro-storage in the smart grid. *Journal of Artificial Intelligence Research (JAIR)*, *42*, 765–813.