

Online Mechanism Design for Electric Vehicle Charging

Enrico H. Gerding*
eg@ecs.soton.ac.uk

David C. Parkes†
parkes@eecs.harvard.edu

Valentin Robu*
vr2@ecs.soton.ac.uk

Alex Rogers*
acr@ecs.soton.ac.uk

Sebastian Stein*
ss2@ecs.soton.ac.uk

Nicholas R. Jennings*
nrj@ecs.soton.ac.uk

*University of Southampton, SO17 1BJ, Southampton, UK

†Harvard University, Cambridge, MA 02138, USA

ABSTRACT

Plug-in hybrid electric vehicles are expected to place a considerable strain on local electricity distribution networks, requiring charging to be coordinated in order to accommodate capacity constraints. We design a novel online auction protocol for this problem, wherein vehicle owners use agents to bid for power and also state time windows in which a vehicle is available for charging. This is a multi-dimensional mechanism design domain, with owners having non-increasing marginal valuations for each subsequent unit of electricity. In our design, we couple a greedy allocation algorithm with the occasional “burning” of allocated power, leaving it unallocated, in order to adjust an allocation and achieve monotonicity and thus truthfulness. We consider two variations: burning at each time step or on-departure. Both mechanisms are evaluated in depth, using data from a real-world trial of electric vehicles in the UK to simulate system dynamics and valuations. The mechanisms provide higher allocative efficiency than a fixed price system, are almost competitive with a standard scheduling heuristic which assumes non-strategic agents, and can sustain a substantially larger number of vehicles at the same per-owner fuel cost saving than a simple random scheme.

Categories and Subject Descriptors

I.2.11 [AI]: Distributed AI - multiagent systems

General Terms

Algorithms, Design, Economics

Keywords

electric vehicle, mechanism design, pricing

1. INTRODUCTION

Promoting the use of electric vehicles (EVs) is a key element in many countries’ initiatives to transition to a low carbon economy [4]. Recent years have seen rapid innovation within the automotive industry [10], with designs such as plug-in hybrid vehicles (PHEVs, which have both an electric motor and an internal combustion engine) and range-extended electric vehicles (which have an electric motor and an on-board generator driven by an internal combustion engine) promising to overcome consumers’ *range anxiety*¹ and thereby increasing mainstream EV use². However, there

¹Fear that a car will run out of electricity in the middle of nowhere.

²The Toyota ‘plug-in’ Prius and the Chevrolet Volt are commercial

are significant concerns within the electricity distribution industries regarding the widespread use of such vehicles, since the high charging rates that these vehicles require (up to three times the maximum current demand of a typical home) could overload local electricity distribution networks at peak times [5]. Indeed, street-level transformers servicing between 10-200 homes may become significant bottlenecks in the widespread adoption of EVs [11].

To address these concerns, electricity distribution companies that are already seeing significant EV use (such as the Pacific Gas and Electric Company in California) have introduced time-of-use pricing plans for electric vehicle charging that attempt to dissuade owners from charging their vehicles at peak times, when the local electricity distribution network is already close to capacity³. While such approaches are easily understood by customers, they fail to fully account for the constraints on the local distribution networks, and they are necessarily static since they require that vehicle owners individually respond to this price signal and adapt their behaviour (i.e., manually changing the time at which they charge their vehicle). Looking further ahead, researchers have also begun to investigate the automatic scheduling of EV charging. Typically, this work allows individual vehicle owners to indicate the times at which the car will be available for charging, allowing automatic scheduling while satisfying the constraints of the distribution network [15, 2]. However, since these approaches separate the scheduling of the charging from the price paid for the electricity (typically assuming a fixed per unit price plan), they are unable to preclude the incentive to misreport (e.g., an owner may indicate an earlier departure time or further travel distances in order to receive preferential charging).

To address the above shortcomings, we turn to the field of *online mechanism design* [12]. Specifically, we focus on mechanisms that are *model-free* (which make no assumptions about future demand and supply of electricity), and that allocate resources as they become available (electricity is *perishable* since installing alternative storage capacity can be very costly). Now, existing mechanisms of this kind assume that the preferences of the agents (representing the vehicle owners) can be described by a single parameter, so-called *single-valued* domains. However, this assumption is not appropriate for our problem, where agents have multi-unit demand with marginal non-increasing valuations for incremental kilowatt hours (kWh) of electricity.⁴ To this end, we extend the state of the art in dynamic mechanism design as follows:

examples of both, which will be on the road in 2011.

³See for example www.pge.com/about/environment/pge/electricvehicles/fuelrates/.

⁴Marginal valuations are non-increasing in our domain because distance and energy usage are uncertain, and therefore the first few units of electricity are more likely to be used, and (in the case of plug-in hybrid electric vehicles) any shortfall can be made up by using the vehicle’s internal combustion engine.

- We develop a formal framework and solution for the EV charging problem, and show that it can be naturally modeled as an online mechanism design problem where agents have multi-unit demand with non-increasing marginal valuations.
- We develop the first model-free online mechanism for perishable goods, where agents have multi-unit demand with decreasing marginal valuations. To ensure truthfulness, we show that this mechanism occasionally requires units to remain unallocated (we say that these units are ‘burned’), even if there is demand for these units. This burning can be done in two ways: at the time of allocation, or on departure of the agent. The latter results in higher allocative efficiency and allocations are easier to compute, but occasionally requires the battery to be discharged which may not always be feasible in practice. Both variants are (weakly) *dominant-strategy incentive compatible* (DSIC), which means that no agent has an incentive to misreport their demand vector and the vehicle availability, regardless of the others’ reports.
- We evaluate our mechanism through numerical simulation of electric vehicle charging using vehicle use data taken from a recent trial of EVs in the UK. In doing so, we show how the agent valuations can be derived from real monetary costs to the vehicle owners, by considering factors such as fuel prices, the distance that the owner expects to travel, and the energy efficiency of the vehicle. Experiments conducted in this realistic setting show that the mechanism with on-departure burning is highly scalable (it can handle hundreds of agents), and both variants outperform any fixed price mechanism for this problem in terms of allocative efficiency, while performing only slightly worse than a well known scheduling heuristic, which assumes non-strategic agents.

Throughout this paper, we focus on measuring allocative efficiency rather than seller profit, since our main design goal is to assure that the capacity of the distribution network is not exceeded, and that agents that need electricity most are allocated, rather than on maximizing profits.

2. RELATED WORK

Online mechanism design is an important topic in the multi-agent and economics literature and there are several lines of research in this field. One of these aims to develop online variants of Vickrey-Clarke-Groves (VCG) mechanisms [13, 7]. While these frameworks are quite general, their focus is on (a slight strengthening of) Bayesian-Nash incentive compatibility, whereas in this paper we focus on the stronger concept of DSIC. Moreover, these works rely on a model of future availability, as well as future supply (e.g., Parkes and Singh [13] use an MDP-type framework for predicting future arrivals), while the mechanism proposed here is model-free. Such models require fewer assumptions, and make computing allocations more tractable than VCG-like approaches.

Model-free settings are considered by both Hajiaghayi et al. [8] and Porter [14], who study the problem of online scheduling of a single, re-usable resource over a finite time period. They characterise truthful allocation rules for this setting and derive lower bound competitive ratios. A limitation of this work [12, 8, 14] is that they consider single-valued domains and, as we show, these existing approaches are no longer incentive compatible for our setting where agents’ preferences are described by a vector of values.

Another related direction of work concerns designing truthful multi-unit demand mechanisms for static settings. A seminal result in this area is the sufficient characterisation of DSIC in terms of weak monotonicity (WMON) [1]. Although this work is relevant

to our model (we briefly discuss the relationship between our mechanism and WMON in Section 4.3), it does not propose any specific mechanism, and, more importantly, existing results do not immediately apply to online domains where agents arrive over time and report their arrival and departure times, as well as their demand.

A different approach for dynamic problems is proposed by Juda and Parkes [9]. They consider a mechanism in which agents are allocated options (a right to buy) for the goods, instead of the goods themselves, and agents can choose whether or not to exercise the options when they exit the market. The concept of options would need to be modified to our setting with perishable goods, with power allocated and then burned so that the final allocation reflects only those options that would be allocated. It is not clear how our online burning mechanism maps to their method.

In addition to theoretical results, several applications have been suggested for online mechanisms, including: the allocation of Wi-Fi bandwidth at Starbucks [6], scheduling of jobs on a server [14] and the reservation of display space in online advertising [3]. However, this is the first work that proposes an online mechanism for electric vehicle charging, and we show how our theoretical framework naturally maps into this domain.

3. EV CHARGING MODEL

In this section we present a model for our problem, formally defining it as an online allocation problem.

(Supply) We consider a model with discrete and possibly infinite time steps (e.g., hourly slots) $t \in T$. At each time step, a number of units of electricity are available for vehicle charging as described by the *supply function* $S : T \rightarrow \mathbb{N}_0^+$, where $S(t)$ describes the number of units available at time t . Supply can vary over time due to changes in electricity demand for purposes other than vehicle charging, as well as changeable supply from renewable energy sources, such as wind and solar.

Importantly, we assume that all vehicle batteries are charged at the same rate.⁵ Thus, a unit of electricity corresponds to the total energy consumed for charging a single vehicle in a single time step. Note that, while there are multiple units of supply at each time step (and agents have demand for multiple units), each agent can be allocated *at most* a single unit per time step. These units are allocated using a periodic multi-unit *auction*, one per time step. Units of electricity are *perishable*, meaning that any unallocated units at each time step will be lost.

(Agents and Preferences) Each vehicle owner is represented by an agent. Let $I = \{1, \dots, n\}$ denote the set of all agents. An agent i ’s (true) availability for charging is given by its *arrival time* $a_i \in T$ (i.e., the earliest possible time the vehicle can be plugged in), and *departure time* $d_i \geq a_i, d_i \in T$ (i.e., after which the vehicle is needed by the owner). We will sometimes use $T_i = \{a_i, \dots, d_i\}$ to indicate agent i ’s availability and we say that agent i is *active* in the market during this period. An agent has a positive value for units allocated when the agent is active, and has zero value for any units allocated outside of its active period. Furthermore, agents have preferences which determine their value or utility for a certain number of units of electricity. These preferences can change from one agent to another, and depend on factors such as the *efficiency* of the vehicle, travel distance, uncertainty in usage, battery capacity and local fuel prices. Formally, preferences are described by a valuation vector $\mathbf{v}_i = \langle v_{i,1}, v_{i,2}, \dots, v_{i,m_i} \rangle$, where $v_{i,k}$ denotes the *marginal value* for the k^{th} unit and m_i is the maximum demand from agent i . That is, $v_{i,k} = 0$ for $k > m_i$. We will often use $v_{i,k+1}$, which describes the value for the next unit when an agent already has k units of electricity. Note that the agent is indifferent

⁵We believe that our approach can be extended to address settings with variable charge rates, but leave this for future work.

w.r.t. the precise allocation times, and merely cares about the total number of units received over the entire active period. These components together describe agent i 's type $\theta_i = \langle a_i, d_i, \mathbf{v}_i \rangle$. We let $\theta = \{\theta_1, \dots, \theta_n\}$, and θ_{-i} is the types of all agents except i . We will often use the notation $(\theta_i, \theta_{-i}) = \theta$.

We assume that agents have *non-increasing marginal valuations*, i.e., $v_{i,k} \geq 0$ and $v_{i,k+1} \leq v_{i,k}$. As we will show in Section 5, this assumption is realistic in a setting with plug-in hybrid and range-extended EVs, where the more a vehicle battery is charged, the less it needs to rely on the fuel-consuming internal combustion engine.

(Reports and Mechanism) Importantly, we allow agents the opportunity to misreport their types. Let $\hat{\theta}_i = \{\hat{a}_i, \hat{d}_i, \hat{\mathbf{v}}_i\}$ denote an agent's report.⁶ Given this, a *mechanism* takes the agents' reported (or observed) types as input as they enter the system, and based on these reports determines the allocation of resources, as well as the payments to the agents. Our goal is then to design a mechanism which incentivises truthful reporting. The *decision policy* then specifies an allocation $\pi_i^{(t)}(\hat{\theta}; \mathbf{k}^{(t)})$ at each time point $t \in T$ and for each agent $i \in I$, where $\mathbf{k}^{(t)} = (k_1^{(t)}, \dots, k_n^{(t)})$ denotes the total *endowments* of the agents at time t before the start of the auction at time t . That is:

$$k_i^{(t)} = \sum_{t'=\hat{a}_i}^{t-1} \pi_i^{(t')}(\hat{\theta}_i, \hat{\theta}_{-i} | \mathbf{k}^{(t')}).$$

The policy π is subject to the constraint that units can only be allocated to agents within their reported activation period. In what follows, we will use the abbreviated notation $\pi_i^{(t)}(\hat{\theta})$, leaving any dependence on the current endowments implicit. Furthermore, let $\pi_i(\hat{\theta}_i, \hat{\theta}_{-i}) = \sum_{t=\hat{a}_i}^{\hat{d}_i} \pi_i^{(t)}(\hat{\theta}_i, \hat{\theta}_{-i})$ denote the total number of units allocated to agent i in its (reported) active time period. We will sometimes omit the arguments when this is clear from the context. Furthermore, the *payment policy* specifies a payment function $x_i(\hat{\theta}_i, \hat{\theta}_{-i} | \pi_i)$ for each agent i . Importantly, while allocations occur at each time point $t \in T$ (since units are perishable), payments are calculated at the reported departure time \hat{d}_i (i.e., when the owner physically unplugs the vehicle).

(Limited Misreports) As in [12], we assume that the agents cannot report an *earlier arrival*, nor a *later departure*. Formally, $\hat{a}_i \geq a_i$ and $\hat{d}_i \leq d_i$, and we say such a pair (\hat{a}_i, \hat{d}_i) is *admissible*. This is a valid assumption in our domain because the agent's vehicle has to be physically plugged into the system, and this cannot be done if the vehicle is not available. However, it can still report an earlier departure since the vehicle can be unplugged before the vehicle is truly needed. Similarly, it can delay its effective arrival (i.e., after having arrived, the vehicle owner can delay actually plugging in the vehicle).

(Agent Utility) Given its preferences, an agent's utility by the departure time is given by the valuation for its obtained units of electricity, minus the payments to the mechanism. Formally:

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

⁶In practice, reported arrival and departure correspond to times when the vehicle is physically plugged into, and, respectively, unplugged from the network (which could differ from when the vehicle is truly available), which can typically be observed by the system. This is because we use a greedy-like scheduling approach (see Section 4) which does not require agents to report their types, nor have knowledge of their true types, in advance. Consequently, it is straightforward to apply our approach to settings where agents do not know their exact availability or this changes due to unexpected events.

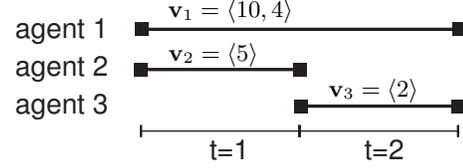


Figure 1: Example showing arrivals, departures, and valuation vectors of 3 agents.

4. THE ONLINE MECHANISM

In this section we consider the problem of designing a model-free mechanism for the above setting. Now, in the case of single-unit demand, a simple greedy mechanism with an appropriate payment policy is DSIC [12]. However, we will show through an example, that this is no longer the case in a multi-unit demand setting that we consider. A greedy allocation is formally defined as follows:

DEFINITION 1 (GREEDY ALLOCATION). *At each step t allocate the $S(t)$ units to the active agents with the highest marginal valuations, $v_{i,k_i^{(t)}+1}$, where ties are broken randomly.*

Consider the example with 2 time steps and 3 agents in Figure 1, showing the agents' arrival, departure and valuations. Suppose furthermore that supply is $S(t) = 1$ at each time step. Greedy would then allocate both units to agent 1, because agent 1 has the highest marginal valuation in both auctions.

Now, consider the question of finding a payment scheme that makes greedy allocation truthful. 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 not be truthful, because it could report $\hat{v}_{1,1} > 5$ 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 2 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 1 unit instead? We argue no more than 2. If $x_1(\hat{\theta} | \pi_1 = 1) = 2 + \epsilon$ (where $\epsilon > 0$), then if the agent's first marginal value was instead $v_{1,1} \in (2, 2 + \epsilon)$, 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 1 unit, not 2, as allocated by the greedy mechanism (its utility for one unit is greater than for two, as $10 - 2 > 10 + 4 - 7$). Hence, online greedy allocation cannot be made truthful.⁷

In order to address this, in our mechanism we extend the Greedy decision policy by allowing the system to occasionally "burn" units of electricity when necessary, in order to maintain incentive compatibility. By burning we mean that this unit is not allocated to any agent, even when there is local demand. We consider two approaches: *immediate burning*, where the decisions to leave a unit unallocated is made at each time step before charging, and *on-departure burning*, where allocated units can be reclaimed by the system when the agent leaves the market (i.e., the corresponding amount of electricity is discharged from the battery on departure).

Each of these approaches has their own advantages and disadvantages. Burning on departure generally requires burning fewer units in some cases, and thus it leads to a higher efficiency. Moreover, the current method we use to determine payments for immediate burning can have a computational cost exponential in the

⁷Formally, this is because the decision policy violates a property called weak monotonicity [1]. In this paper, we omit a detailed discussion of this relationship, due to space restrictions.

number of the agents present, whereas for on-departure burning, the cost of determining payments is linear. However, in terms of the application domain, fast discharging of a vehicle's battery may not be practical.

Note that, for both approaches, the energy that is burnt is not necessarily wasted, but it is simply returned to the grid, to be used for other purposes than electric vehicle charging. For immediate burning, the unallocated electricity units are returned to the grid before it is actually charged by the agent. For the mechanism with on-departure burning, units may be charged first and then rapidly discharged when the agent leaves the market. While this may result in some loss, this is probably negligible w.r.t. the overall amount of electricity allocated.

4.1 The Mechanism

Before we introduce the decision policy, we show how we can compute a set of threshold values, which are used both to calculate the payments and to decide when to burn a unit of electricity. Let $k_{-i,j}^{(t)} = \sum_{t'=a_j}^{t-1} \pi_j^{(t')}(\theta_{-i})$ denote the endowment of an active agent j at start time t , under the allocation we would have *in the absence of agent i* (note that calculating this value requires recomputing allocations without agent i in the market from a_i until the current time t). Then $v_{j,k_{-i,j}^{(t)}+1}$ is the marginal valuation of agent j at time t in the absence of agent i . Given this, we define $v_{-i,t}^{(n)}$ to be the n^{th} highest of such valuations from all active agents $j \neq i$. Then $v_{-i,t}^{(S(t))}$, for supply $S(t)$, is the lowest value that is still allocated a unit at time t , if agent i were not present. Henceforth, we refer to $v_{-i,t}^{(S(t))}$ as the *marginal clearing value* for agent i in period t , and we will often use $v_{-i,t} = v_{-i,t}^{(S(t))}$ for brevity.

Now, let $\mathbf{p}_{-i}^{(t)} = \text{incr}(v_{-i,a_i}, v_{-i,a_i+1}, \dots, v_{-i,t})$ denote agent i 's price vector at time t , where a_i is the reported arrival time of agent i and $\text{incr}(\cdot)$ is an operator which takes a vector of real values as input and returns it in increasing order. In addition, let $\mathbf{p}_{-i} = \mathbf{p}_{-i}^{(d_i)}$ denote the value of this vector at time d_i when agent i leaves the market.

Intuitively, in any round t , the price $p_{-i,k}^t$ that agent i is charged for the k -th unit is the minimum valuation the agent could report for the k -th unit and win it by time t , given the greedy allocation policy with burning described below. Given this, the decision and payment policies of our mechanism are as follows.

- **Decision Policy** The decision consists of two stages.

Stage 1 At each time point t , *pre-allocate* using Greedy (see Definition 1).

Stage 2 We consider two variations in terms of when to decide to burn pre-allocated units:

- **Immediate Burning.** Burn a unit whenever:

$$v_{i,k_i^{(t)}+1} < p_{-i,k_i^{(t)}+1}^{(t)}$$

- **On-Departure Burning.** This type of burning occurs on reported departure. For each departing agent, burn any unit $k \leq \pi_i$ where $v_{i,k} < p_{-i,k}$.

- **Payment Policy** Payment occurs on reported departure. Given that π_i units are allocated to agent i at time $t = \hat{d}_i$, the payment collected from i is:

$$x_i(\hat{\theta}_i, \hat{\theta}_{-i}|\pi_i) = \sum_{k=1}^{\pi_i} p_{-i,k} \quad (2)$$

Burning occurs whenever the marginal value for an additional unit is smaller than the marginal payment for that unit. Thus these values are effectively agent-specific threshold values, below which no

	agent 1: $T_1 = \{1, 2, 3\}$ $\mathbf{v}_1 = \langle 10, 4 \rangle$	agent 2: $T_2 = \{1\}$ $\mathbf{v}_2 = \langle 5 \rangle$	agent 3: $T_3 = \{2, 3\}$ $\mathbf{v}_3 = \langle 2 \rangle$
$t = 1$	$k_1^{(1)} = 0$ $v_{-1,1} = 5$ $\mathbf{p}_{-1}^{(1)} = \langle 5 \rangle$ $\pi_1^{(1)} = 1$	$k_2^{(1)} = 0$ $v_{-2,1} = 10$ $\mathbf{p}_{-2}^{(1)} = \langle 10 \rangle$ $\pi_2^{(1)} = 0$	
$t = 2$	$k_1^{(2)} = 1$ $v_{-1,2} = 2$ $\mathbf{p}_{-1}^{(2)} = \langle 2, 5 \rangle$ $\pi_1^{(2)} = 0$ (IM) $\pi_1^{(2)} = 1$ (OD)		$k_3^{(2)} = 0$ $v_{-3,2} = 4$ $\mathbf{p}_{-3}^{(2)} = \langle 4 \rangle$ $\pi_3^{(2)} = 0$
$t = 3$ IM	$k_1^{(3)} = 1$ $v_{-1,3} = 0$ $\mathbf{p}_{-1}^{(3)} = \langle 0, 2, 5 \rangle$ $\pi_1^{(3)} = 1$		$k_3^{(3)} = 0$ $v_{-3,3} = 4$ $\mathbf{p}_{-3}^{(3)} = \langle 4, 4 \rangle$ $\pi_3^{(3)} = 0$
$t = 3$ OD	$k_1^{(3)} = 2$ $v_{-1,3} = 0$ $\mathbf{p}_{-1}^{(3)} = \langle 0, 2, 5 \rangle$ $\pi_1^{(3)} = 0$		$k_3^{(3)} = 0$ $v_{-3,3} = 0$ $\mathbf{p}_{-3}^{(3)} = \langle 0, 4 \rangle$ $\pi_3^{(3)} = 1$

Table 1: Example run of the mechanism with 3 agents and 3 time periods for immediate (IM) and on-departure (OD) burning. Grey cells indicate different values for IM and OD burning.

unit is allocated to that agent. Moreover, it is important to note that the mechanism used for computing the prices mirrors the actual allocation mechanism. So, for example, if immediate burning is used in the decision policy, then for each agent i and for all times t , the values of the $\mathbf{p}_{-i}^{(t)}$ vector are computed by re-running the market, in the absence of agent i using immediate burning, based on the reports of the other agents. Conversely, if on-departure burning is used for the decision policy, the same mechanism should be used in computing the \mathbf{p}_{-i} prices.

4.2 Example

To demonstrate how the mechanism works, we extend the previous example 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 3 units to be allocated in total. Table 1 shows the endowments $k_i^{(t)}$, the marginal clearing values $v_{-i,t}$, the $\mathbf{p}_{-i}^{(t)}$ vectors, and the allocation decisions $\pi_i^{(t)}$ at different time periods.

We start by considering the allocations and payment using *immediate burning*. At time $t = 1$, Stage 1 of the mechanism allocates the unit to agent 1, and since $v_{1,1} = 10 \geq p_{-1,1}^{(1)} = 5$, this unit is not burnt in the second stage. At time $t = 2$, the unit again gets pre-allocated to agent 1 since $v_{1,2} = 4 > v_{3,1} = 2$. However, the marginal clearing value $v_{-1,2}$ is inserted at the beginning of the $\mathbf{p}_{-1}^{(2)}$ vector, and as a result $v_{1,2} = 4 < p_{-1,2}^{(2)} = 5$. Consequently, this unit gets burnt and is allocated to neither of the agents. At time $t = 3$, therefore, the marginal value of agent 1 is still 4 (since its endowment is unchanged), and this value is added to agent 3's marginal clearing values. To calculate the marginal clearing value of agent 1, recall that the decision 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 marginal value of this agent is 0. Thus, the value of 0 is inserted in the $\mathbf{p}_{-1}^{(3)}$ vector. At $t = 3$, since agent 1 still has the highest marginal value, it is again pre-allocated the unit. However, now $v_{1,2} = 4 \geq p_{-1,2}^{(3)} = 2$, and therefore the unit is not burnt. So, in case of immediate burning, 2 out of 3 units are allocated to agent

1, and that agent pays $\mathbf{p}_{-1,1}^{(3)} + \mathbf{p}_{-1,2}^{(3)} = 2$.

Now consider the same setting but with *on-departure burning*. The first two time steps are as before, except that there is no burning 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 marginal value of agent 1 at $t = 3$ is equal to $v_{1,3} = 0$. Therefore, the unit is allocated to agent 3, and the payment for this unit is $p_{-3,1} = 0$. The vector $\mathbf{p}_{-1}^{(3)}$ remains unchanged compared to the immediate burning case. At this point, there is no longer a need to burn one of the units of agent 1, since it has received $k = 2$ units, the same allocation as with immediate burning, and note that $v_{1,2} > p_{-1,2}$.

Still, it is possible to construct examples where, both with on-departure and immediate burning, half of the units need to be burnt. Furthermore, note that this unit cannot go to agent 3, because payment would have been $p_{-3,1}^{(3)} = 4$, which would result in a negative utility for agent 3.

4.3 Properties

In this section we prove that the above mechanism is DSIC. 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 function, and given any admissible pair $\langle \hat{a}_i, \hat{d}_i \rangle$, if a decision policy is a threshold policy, then the mechanism is DSIC with respect to the valuations (Lemma 1). (ii) We show that our decision policy is a threshold policy (Lemma 2). (iii) Finally, we show that, if agents truthfully report their valuations, reporting $\hat{a}_i = a_i, \hat{d}_i = d_i$ is a weakly dominant strategy (Lemma 3). These results are combined in Theorem 1 to show that our policy is DSIC.

DEFINITION 2 (THRESHOLD POLICY). *A decision policy π is a threshold policy if, for a given agent i with fixed $\langle \hat{a}_i, \hat{d}_i \rangle$ 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 $\hat{v}_{i,k} \geq \tau_k$.*

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 .⁸

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 $\hat{\mathbf{v}}_i$. Note that 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 $\langle \hat{a}, \hat{d} \rangle$ and $\hat{\theta}_{-i}$, if π is a threshold policy coupled with a payment policy:*

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

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

⁸A threshold policy satisfies weak-monotonicity (WMON) [1], and is therefore sufficient for truthfulness 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. We show that our decision policy has the threshold property, and thus the WMON, and that it also handles misreports of arrivals and departures.

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

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

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 2 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:

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

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:

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

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. \square

Note that Greedy (as per Definition 1) is not a threshold policy. To see this, consider the example from Figure 1. As we saw earlier, Greedy allocates 2 units to agent 1, and the required threshold τ_2 for winning the second unit is 2 (below which Greedy would allocate 1 unit). However, if agent 1 had valuation $\mathbf{v}_1 = \langle 4, 4 \rangle$, Greedy would allocate only 1 unit, even though $v_2 > \tau_2$, which conflicts with the requirement of a threshold policy.

The next lemma shows that the threshold condition holds if we include burning, and if we set the threshold values to $\tau_k = p_{-i,k}$.

LEMMA 2. *Given non-increasing marginal valuations, the decision policy π in Section 4.1 is (for either burning policy) a threshold policy where $\tau_k = p_{-i,k}$.*

PROOF. First, from the definition of vector $\mathbf{p}_{-i}^{(t)}$ and \mathbf{p}_{-i} from Section 4.1, the values of $\mathbf{p}_{-i}^{(t)}$ are independent of the reports $\hat{\mathbf{v}}_i$ made by agent i . This is because each of its component values $v_{-i,a_i}, \dots, v_{-i,t}$ are computed based only on the reports of the other agents, by first removing agent i from the market.

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 less than k units to agent i .

Part 1: Let $v_{i,k} \geq p_{-i,k}$. Suppose that agent i has the same marginal values, $v_{i,k}$, for the first k units (i.e., $v_{i,1} = v_{i,2} = \dots = v_{i,k}$), then it will win exactly those auctions where $v_{i,k} \geq v_{-i,t}$, $t \in T_i$ in Stage 1 of the mechanism (ignoring tie breaking). Note that even when, by winning an auction, agent i displaces the losing marginal value to a future auction, since this value is less or equal to $v_{i,k}$, it will not affect the future auctions for agent i since it will still outbid that agent in the next auction. Now, because $p_{-i,j} \leq p_{-i,k}$ for $j \leq k$ (by definition), there must be at least k auctions where $p_{-i,k} \geq v_{-i,t}$ in the period $t \in T$, and since $v_{i,k} \geq p_{-i,k}$, agent i wins at least k auctions in Stage 1.

Furthermore, each time an auction is won, the clearing values appear as one of the j first elements of the \mathbf{p}_{-i}^t vector, where j is the number of auctions so far (since these are the auctions with the lowest clearing values, and the clearing values are ordered ascendingly). Because agent i wins an auction in Stage 1 if and only if $v_{i,k} \geq v_{-i,t}$, it follows that $v_{i,k} = v_{i,j} \geq p_{-i,j}$ whenever it wins an auction in Stage 1. Therefore, there is no burning 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 decision policy will allocate *at least* k units to agent i .

Part 2: Let $v_{i,k} < p_{-i,k}$. First consider the *on-departure burning* case. As per the definition of Stage 2 of the mechanism, unit k is burnt. However, we still need to show that any units $j > k$ are burnt as well. Since $p_{-i,j} \geq p_{-i,k}$ and $v_{-i,j} \leq v_{-i,k}$ for all $j > k$, it follows that $v_{i,j} < p_{-i,j}$ for all $j > k$. Therefore even if Stage 1 allocates k or more units, these will be burnt in Stage 2, and thus strictly less than k units remain.

Now consider the *immediate burning* case. Note that $p_{-i,k} \leq p_{-i,k}^{(t)}$ for $(a_i + k - 1) \leq t \leq d_i$. That is, threshold values can only decrease over time. Thus it follows that $v_{-i,k} < p_{-i,k}^{(t)}$ 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. Because $v_{-i,k} < p_{-i,k}^{(t_k)}$, this unit will always be burnt in Stage 2 of the decision policy. Therefore, the final result is an allocation of strictly less than k units. \square

By setting $\tau_k = p_{-i,k}$, the payment function in Equation 2 corresponds to the payment function in Lemma 1. Therefore the proposed mechanism is shown to be DSIC in valuations. We now complete the proof by showing that truthful reporting of the arrival and departure times are also DSIC (given limited misreports), given 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 pair of arrival/departure reports (\hat{a}_i, \hat{d}_i) , then it is a dominant strategy to report $\hat{a}_i = a_i$ and $\hat{d}_i = d_i$.*

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

$\pi_i(\hat{\theta}_i) = \pi_i(\theta_i)$: Misreporting by agent i has no affect on the marginal clearing values $v_{-i,t}$, but can only decrease the size of the \mathbf{p}_{-i} vector. In particular, due to limited misreports we have $\hat{a}_i \geq a_i$ and $\hat{d}_i \leq d_i$, and thus $\mathbf{p}_{-i}^{(\hat{a}_i, \hat{d}_i)}$ contains a *subset* of the elements from $\mathbf{p}_{-i}^{(a_i, d_i)}$. As these vectors are by definition increasingly ordered, it follows that $p_{-i,j}^{(\hat{a}_i, \hat{d}_i)} \geq p_{-i,j}^{(a_i, d_i)}, \forall j \leq (\hat{d}_i - \hat{a}_i + 1)$. Since the payment consists of the first $k_i = \hat{k}_i$ elements, this can only increase by misreporting.

$\pi_i(\hat{\theta}_i) \neq \pi_i(\theta_i)$: First, we show that $\pi_i(\hat{\theta}_i) > \pi_i(\theta_i)$ could never occur. Since the marginal clearing values remain the same, but the number of auctions in which the agent participates decreases by misreporting, Stage 1 of the mechanism can only allocate fewer or equal items. Furthermore, since $p_{-i,j}^{(\hat{a}_i, \hat{d}_i)} \geq p_{-i,j}^{(a_i, d_i)}$, the possibility of burning 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 $\sum_{j=1}^{\pi_i(\hat{\theta}_i)} p_{-i,j}^{(\hat{a}_i, \hat{d}_i)} \leq \sum_{j=1}^{\pi_i(\hat{\theta}_i)} p_{-i,j}^{(a_i, d_i)}$ (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. \square

We are now ready to present the main theoretical result:

THEOREM 1. *Given non-increasing marginal valuations and limited misreports, Greedy with on-departure and immediate burning and with payment function according to Equation 2 are DSIC.*

PROOF. The proof of this theorem follows directly from the above lemmas. Lemmas 1 and 2 show that, for any pair of arrival/departure (mis)-reports (\hat{a}_i, \hat{d}_i) the decision policy is truthful in terms of the valuation vector \mathbf{v}_i , given an appropriate payment policy. Furthermore, the payments in Equation 2 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 by misreporting arrivals/departures. \square

5. EXPERIMENTAL EVALUATION

In this section, we evaluate our proposed mechanism empirically. In doing so, we seek to answer a number of pertinent questions. First, since our greedy approach does not generally find the optimal allocation, we are interested in how close it comes to this in realistic settings. Second, we investigate the extent to which unit burning occurs in practice (i.e., how often units of electricity need to be burned by our decision policies, in order to ensure truthfulness). This is critical, as it may negatively affect efficiency. Finally, we compare our mechanism to a range of simpler truthful mechanisms that employ fixed pricing, as well as to a well-known online scheduling approach. These serve as benchmarks for our mechanism — fixed pricing is a common mechanism for selling goods in a wide range of settings, while the scheduling approach highlights what a non-truthful mechanism could achieve.

5.1 Experimental Setup

Our experimental setup is based on data collected during the first large-scale UK trial of EVs. In December 2009, 25 EVs were provided to members of the public as part of the CABLED (Coventry And Birmingham Low Emissions Demonstration) project.⁹ The aim of this trial was to investigate real-world usage patterns of EVs. To this end, they were equipped with GPS and data loggers to record comprehensive usage information, such as trip durations and distances, home charging patterns and energy consumption.

We use the data published by this project for the first quarter of 2010 to realistically simulate typical behaviour patterns. More specifically, in each of our experiments, we simulate a single 24 hour day, where charging periods are divided into hourly time intervals. For the purpose of the experiments, a simulated day starts at 15:00, as vehicle owners begin to arrive back from work. To determine the arrival time of each agent, we randomly draw samples from the home charging start times reported by the project. These are highest after 18:00 and then quickly drop off during the night. Likewise, to simulate departures, we sample from data recording journey start times.

In order to simulate realistic marginal valuation vectors for the agents, we combine data from the project about journey distances with a principled approach for calculating the expected economic benefit of vehicle charging. In particular, we can calculate the expected utility of a given amount of charge (in kWh), c_e , given a price of fuel (in £/litre), p_p , an internal combustion engine efficiency (in miles/litre), e_p , an electric efficiency (in miles/kWh), e_e , and a probability density function, $p(m)$, that describes the distance to be driven the next day:

$$\mathbb{E}(u(c_e)) = \int_0^\infty \frac{p_p}{e_p} \cdot m \cdot p(m) dm - \int_{c_e \cdot e_e}^\infty \frac{p_p}{e_p} \cdot m \cdot p(m) dm, \quad (3)$$

⁹See <http://cabled.org.uk/>.

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 . Given this, and a charging rate (in kW), r_e , it is straight-forward to calculate the marginal valuation of the k th hour of charging time: $v_k = \mathbb{E}(u(k \cdot r_e)) - \mathbb{E}(u((k-1) \cdot r_e))$.

To generate a variety of marginal valuations, we note that e_e and e_p depend on the specific make and type of the EV and thus vary between households, while $p(m)$ depends on the driving behaviour of the car owner. We draw e_e uniformly at random from 2 – 4 miles/kWh and e_p is drawn from 9 – 18 miles/litre. Furthermore, we create $p(m)$ from daily driving distances presented in the CABLED report. These distances are typically short, with a daily mean of 23 miles, but the distribution has a long tail with a maximum of 101 miles. Next, we draw the capacity of a car battery from 15 – 25 kWh and set the charging rate to 3 kW. These and earlier specifications are all based on the Chevrolet Volt, the first mass-produced range-extended EV to be on the road in 2011. However, we include some variance to account for other vehicle types.

Finally, to derive the supply function S , we consider a realistic *neighbourhood-based* supply function using the average energy consumption of a UK household over time.¹⁰ In this setting, the total energy available for charging depends on the number of households in the neighbourhood and the constraints of the local transformer. Hence, available supply during the night is significantly higher than during the day. Furthermore, we tested a range of other supply functions and valuation distributions, where we observed the same general trends as discussed in the remainder of this section. However, we omit the details here for brevity.

5.2 Benchmark Mechanisms

In addition to the two decision policies developed within this paper — Greedy with Immediate Burning (*Immediate*) and Greedy with On-Departure Burning (*On-Departure*) — we benchmark the following strategies that have been widely applied in similar settings:

Fixed Price allocates units to those agents that value them higher than a fixed price p . The price they pay for this unit is p . When demand is greater than supply, units are allocated randomly between all agents with a sufficiently high valuation. 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, we test all possible values (in steps of £0.01) and select the p that achieves the highest average efficiency (over 1000 trials) for a given setting. Thus, when showing the results of *Fixed Price*, this constitutes an upper bound of what could be achieved with this mechanism. We use the special case $p = 0$ as a baseline benchmark and denote this as *Random*.

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, 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 .

Optimal allocates units to agents to maximise the overall allocation efficiency, assuming complete knowledge of future arrivals and supply. Clearly, this mechanism is not practical 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.

Having described the valuation calculation, the experimental setting, and the benchmarks, we now describe our results.

¹⁰We use the average evaluated during a work day in winter, available at <http://www.elexon.co.uk/>.

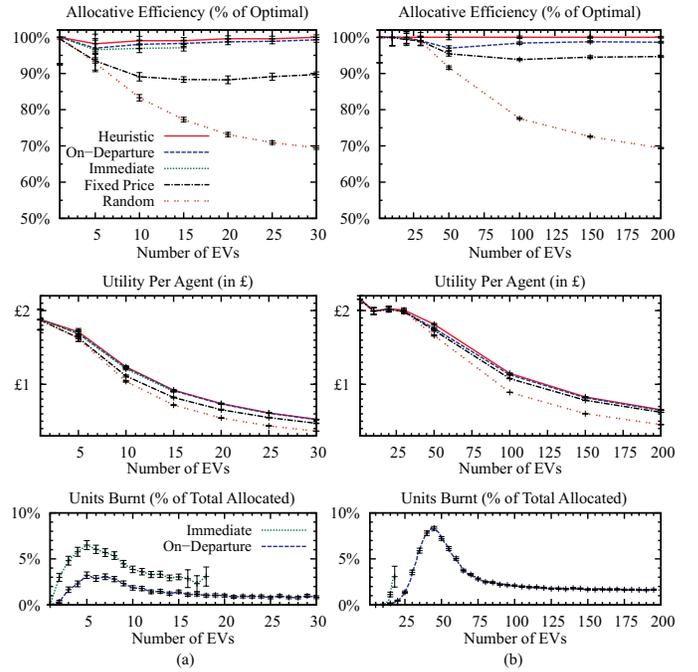


Figure 2: Results for a small neighbourhood with 30 houses (a) and a large one with 200 houses (b).

5.3 Results

For our experiments, we consider two possible neighbourhood sizes — one with 30 households and one with 200 households. In these settings, the capacity of the local transformer is constrained, so that only a couple of cars can charge at the same time in the 30 household case and up to 16 with 200 households. We choose such highly constrained settings here, because they are intrinsically more challenging and interesting than settings where all cars can be fully charged overnight. Across the experiments, we vary the number of these households that own an EV. Note here that we only show results for *Immediate* burning up to 15 agents, because our current implementation of this is computationally expensive. This is because the vector of marginal clearing values $\mathbf{p}_{-i}^{(t)}$ at time t depends on which units are burned in i ’s absence (and as this vector is used to determine when burning takes place, it recursively depends on the corresponding vectors of all agents that are allocated in i ’s absence). Thus, we may potentially need to evaluate all subsets of agents, which grows exponentially with n . Although it may be possible to prune the search space efficiently in practice, we leave these computational aspects to future work. It is interesting that this does not apply to *On-Departure* burning, because here burning does not influence the agents’ marginal clearing values.

The results for both settings are given in Figure 2. First, the top row shows the average¹¹ efficiency, normalised to the performance of *Optimal* (when there are more than 30 EV owners, *Optimal* becomes intractable and so we normalise results to the performance of *Heuristic* in those cases as a close approximation). Here, we note that our two burning policies consistently outperform (or match) all other truthful benchmarks. The improvement compared to *Random* is particularly pronounced, but our approach still achieves a significant improvement over the *Fixed Price* mechanism. For small neighbourhoods, this is almost 10%, while in larger neighbour-

¹¹All results are averaged over 1000 trials. We plot 95% confidence intervals, and significant differences reported are at $t < 0.05$ level.

hoods, it is up to 5%. This is a promising result, because setting the optimal price for the fixed price strategy requires knowledge about the distributions of agents types, but our approach makes no such assumptions.

This improvement is due the ability of our mechanism to allocate the agents with the highest marginal valuations, while *Fixed Price* randomises over those that meet its price. Our approach is also responsive to changes in demand over time, consistently allocating units even when the highest valuations are low. In contrast, *Fixed Price* must be tuned to operate at any particular balance of supply and demand. Thus, it does not allocate when its price is unmet. It performs better in the larger setting because it is more likely that at least some of the agents meet the fixed price in this case.

Next, our mechanism also performs close to the *Optimal* and *Heuristic*, consistently achieving 95% or better, which indicates that our greedy approach performs well in realistic settings even without having access to complete information (such as departure times or even future arrivals). The lowest relative efficiency to the optimal is achieved when there are few EVs (about 20% of the neighbourhood). Here, scheduling constraints are most critical, as it may sometimes be optimal to prioritise an agent with lower valuations over one with higher valuations, but a longer deadline. This becomes less critical when there are more agents, as there are typically sufficiently many with high valuations. Finally, we see that *Immediate* burning achieves a slightly lower average efficiency than *On-Departure*. This is due to higher levels of burning, but the difference is small (and, in fact, not statistically significant).

In the second row of Figure 2, the average utility of each EV owner's allocation (not including the payments to the mechanism) is shown. This corresponds directly to the fuel costs that a single EV owner saves by using electricity instead of fuel. Initially, this is high (around £2), as there is little competition, but starts dropping as more EV 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 incentive-compatible mechanisms. For example, to save at least £1 per agent in the small neighbourhood, *Random* can support up to 10 EV owners, while *Immediate* and *On-Departure* achieve the same threshold for up to 14 EV owners (a 40% improvement). In the large neighbourhood, our mechanism can support around 60 additional vehicles in some cases (to achieve a £0.65 threshold).

Finally, the last row shows the average number of units that are burned by our two decision policies, as a percentage of the overall (tentatively) allocated units. Again, due to computational limitations, full results for the *Immediate* burning policy are only shown up to 15 agents. For up to 18 agents, results from only 100 trials are shown (resulting in larger confidence intervals). *On-Departure* burning clearly burns significantly fewer units than *Immediate*, as the latter sometimes unnecessarily burns units. There is also a clear maximum in the number of burned units when around 20% of households are EV owners. This is because there is a significant amount of competition, with many agents that have similar marginal valuations, and this induces burning. However, when the number of agents rises further, burning drops again. This is because agents are increasingly less likely to be allocated more than a single unit in these very competitive settings and so there is no need for burning. It should be noted that burning is generally low (for *On-Departure* burning), with typically only 1-2% of allocated units being burned (and always less than 10%).

6. CONCLUSIONS

This paper proposes a novel online allocation mechanism for a problem that is of great practical interest for the smart grid community, that of integrating EVs into the electricity grid. Our contri-

bution to existing literature is two-fold. On the theoretical side, we extend model-free, online mechanism design with perishable goods to handle multi-unit demand with decreasing marginal valuations.

On the practical side, we empirically evaluate our mechanism in a real-world setting, and showed that the proposed mechanism is highly robust, and achieves better allocative efficiency than any fixed-price benchmark, while only being slightly suboptimal w.r.t. an established cooperative scheduling heuristic.

For future work we plan to look at several issues. First, in this paper we assumed all EVs have a uniform charging rate, but in the future we plan to extend the allocation model to deal with heterogeneous maximal charging rates (corresponding to different types of EVs). Second, it would be interesting to compare the performance of the model-free online mechanism proposed in this paper to a model-based approach, such as the one in [13]. Finally, this paper looked at performance in terms of a realistic application scenario, but we also plan to study the worst-case bounds on allocative efficiency and number of items our mechanism burns in future work.

7. REFERENCES

- [1] S. Bikhchandani, S. Chatterji, R. Lavi, A. Mu'alem, N. Nisan, and A. Sen. Weak monotonicity characterizes deterministic dominant-strategy implementation. *Econometrica*, 74(4):1109–1132, 2006.
- [2] K. Clement-Nyns, E. Haesen, and J. Driesen. The impact of charging plug-in hybrid electric vehicles on a residential distribution grid. *IEEE Transactions on Power Systems*, 25(1):371–380, 2010.
- [3] F. Constantin, J. Feldman, S. Muthukrishnan, and M. Pal. An online mechanism for ad slot reservations with cancellations. In *Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA'09)*, pages 1265–1274, 2009.
- [4] Department of Energy and Climate Change. *The UK low carbon transition plan: National strategy for climate and energy*. HM Government, 2009.
- [5] P. Fairley. Speed bumps ahead for electric-vehicle charging. *IEEE Spectrum*, Jan. 2010. Available at <http://spectrum.ieee.org/green-tech/advanced-cars/speed-bumps-ahead-for-electricvehicle-charging/>.
- [6] E. Friedman and D.C. Parkes. Pricing WiFi at Starbucks— Issues in online mechanism design. In *Proc. of the 4th ACM Conf. on Electronic Commerce*, pages 240–241, 2003.
- [7] A. Gershkov and B. Moldovanu. Efficient sequential assignment with incomplete information. *Games and Economic Behavior*, 68(1):144–154, 2010.
- [8] M. Hajiaghayi, R. Kleinberg, M. Mahdian, and D.C. Parkes. Online auctions with re-usable goods. In *Proc. of the 6th ACM Conf. on Electronic Commerce (EC'05)*, pages 165–174, 2005.
- [9] A.I. Juda and D.C. Parkes. An options-based solution to the sequential auction problem. *Artificial Intelligence*, 173:876–899, 2009.
- [10] W.J. Mitchell, C.E. Borroni-Bird, and L.D. Burns. *Reinventing the automobile: Personal urban mobility for the 21st century*. MIT Press, 2010.
- [11] Royal Academy of Engineering. *Electric Vehicles: Charged with potential*. Royal Academy of Engineering, 2010.
- [12] D.C. Parkes. Online mechanisms. In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, pages 411–439, 2007.
- [13] D.C. Parkes and S. Singh. An MDP-Based approach to Online Mechanism Design. In *Proc. of NIPS'03*, 2003.
- [14] R. Porter. Mechanism design for online real-time scheduling. In *Proc. the 5th ACM Conf. on Electronic Commerce (EC'04)*, pages 61–70, 2004.
- [15] S. Vandael, N. Boucke, T. Holvoet, and G. Deconinck. Decentralized demand side management of plug-in hybrid vehicles in a smart grid. In *Proc. of 1st Int. Workshop on Agent Technologies for Energy Systems*, pages 67–74, 2010.