

Two-Sided Online Markets for Electric Vehicle Charging

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

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

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

Dengji Zhao†
dzhao@scm.uws.edu.au

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

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

†Intelligent Systems Lab, University of Western Sydney, Australia

ABSTRACT

With the growing popularity of electric vehicles (EVs), the number of public charging stations is increasing rapidly, allowing drivers to charge their cars while parked away from home or en-route to their destination. However, as a full charge can take a significant amount of time, drivers may face queues and uncertainty over availability of charging facilities at different stations and times. In this paper, we address this problem by proposing a novel, two-sided market for advance reservations, in which agents, representing EV owners, report their preferences for time slots and charging locations, while charging stations report their availability and costs. In our model, both parties are rational, profit-maximising entities, and buyers enter the market dynamically over time. Given this, we apply techniques from online mechanism design to develop a pricing mechanism which is truthful on the buyer side (i.e., drivers have no incentive to misreport their preferences or to delay their reservations). For the seller side, we adapt three well-known pricing mechanisms and compare them both theoretically and empirically. Using realistic simulations, we demonstrate that two of our proposed mechanisms consistently achieve a high efficiency (90–95% of optimal), while offering a trade-off between stability and budget balance. Surprisingly, the third mechanism, a common payment mechanism that is truthful in simpler settings, achieves a significantly lower efficiency and runs a high deficit.

Categories and Subject Descriptors

I.2.11 [AI]: Distributed AI—Multiagent systems

Keywords

Electric vehicles; mechanism design; two-sided markets

1. INTRODUCTION

Recent years have seen increasing interest in electric vehicles (EVs) as a key technology for achieving efficient transportation with low carbon emissions [1]. However, large-scale use of EVs will raise a host of new challenges for electricity distribution networks [5, 14]. More specifically, electric vehicles are high electricity consumers and, moreover, charging an electric vehicle takes considerably more time than fueling a petrol-powered vehicle.¹ Thus, the problem of efficiently scheduling the charging of a large number of EVs at multiple charging stations will become increasingly pressing and challenging, especially as both electric vehicle owners and charging stations can be seen as self-interested parties, interested in minimising their costs, or maximising their profits, respectively.

¹Fully charging an EV takes a minimum of half an hour, even with the fastest available chargers on the market today.

To address this problem, we present the first system where EVs are matched to charging stations in a two-sided online market. In this system, agents representing EV drivers enter the marketplace dynamically over time, at which point they place bids for *advance reservations* on behalf of their owners. Through these bids, agents express their preferences for different time slots and charging stations. At the same time, charging stations can offer available charging units through the reservation system, and report their minimum prices for different time slots. The system then allocates buyers to these advance time slots *online*, i.e., as they enter the marketplace.

The proposed system is very general and we show how it can be used in two specific real-world scenarios: (1) *park 'n charge*, where the EV is charged while parked at a convenient location away from home and (2) *en-route charging*, where the EV requires charging on the way to a destination. In these scenarios, we envision that the agent is integrated with an automated advice interface and a route planner, which enables the agent to trade off price, availability and distance, and automatically re-routes the user to the relevant stations. Online systems exhibiting some of these features are already beginning to emerge. For example, *Google Maps*² provides interactive directions, allowing drivers to make informed choices between multiple routes based on distance, estimated time of driving or fuel costs. In terms of EV charging, companies such as *PlugShare*³ and *ChargePoint*⁴ provide interactive maps of available EV charging points in the US and Canada (including some reservation facilities).

Our work is closely related to combinatorial exchanges, where buyers and sellers are matched based on their (combinatorial) preferences. These are mainly concerned with finding allocations and payments that incentivise truthful bidding, while satisfying other properties such as individual rationality (buyers and sellers make no loss from participating) and budget balance (the system or auctioneer makes no loss). A seminal result in this field is by McAfee [9], who proposes a payment scheme to achieve truthfulness in two-sided markets with identical goods. Gonen *et al.* [7] extend this to combinatorial two-sided auction markets, and propose a procedure called generalised trade reduction in order to ensure several economic properties including truthfulness. However, this and similar work assumes a static setting where all buyers and sellers are all present in the market at the same time. In our setting, on the other hand, buyers enter the system dynamically and need to be allocated without knowing future demand.

Incentivising truthful behaviour in settings with dynamic market entry is studied in the field of *online mechanism design* (see [11, Ch.16] for a survey). Existing work, such as [12, 6], largely con-

²<http://maps.google.com/>

³<http://www.plugshare.com/>

⁴<http://www.chargepoint.net/>

siders one-sided markets (e.g., with one seller and many buyers). In contrast, we look at two-sided markets with multiple buyers (the EVs) and multiple sellers (the charging stations). Online two-sided markets are studied in [3, 4], but all buyers and sellers are assumed to trade a single unit of the same commodity. Our setting is much more complex, since buyers have different preferences for different sellers and time slots, and sellers can have multiple time slots and multiple units per time slot. As we will show, this added complexity has significant implications for the properties of the market.

More recently, some research has attempted to address the EV charging problem from a cooperative scheduling perspective. In this vein, [13, 8] consider constructing an online charging schedule which takes into account spatial and temporal constraints. While some of these approaches use similar concepts to our work, such as placing advance reservations while en-route (e.g., [8]), they rely on a centralised scheduler and fully cooperative agents. In contrast, we assume that agents are self-interested and can strategise by misreporting their preferences. Furthermore, we propose a decentralised marketplace where agents perform much of the computation, such as the routing and computing the EV’s energy requirements. A small number of papers have studied the EV scheduling problem considering strategic, self-interested agents [6, 16, 15]. However, these study a one-sided setting with a single, fixed charging point.

Against this background, this paper makes the following contributions to the existing state of the art:

- We introduce the first two-sided market architecture for matching EV owners to charging stations using advance reservations.
- For this system, we develop a payment mechanism that is truthful and individually rational on the buyer side. On the seller side, we outline a number of payment mechanisms and explore their theoretical properties. As part of this, we present an impossibility result which shows that no payment can always be truthful for sellers when a greedy allocation rule is used.
- We show how our reservation system can be applied to two realistic scenarios, and we analyse the equilibrium behaviour of agents in these scenarios using extensive simulations. We demonstrate that two of our proposed payment mechanisms induce a high allocative efficiency (around 90–95% of the optimal), and we show that one of these achieves a higher stability at the expense of running a small deficit. Surprisingly, we find that a well-known payment mechanism that is truthful for sellers in simpler settings performs poorly, in terms of both efficiency and deficit.

The remainder of the paper is organised as follows. In Section 2 we first present our system model. In Section 3 we describe the allocation mechanism and several payment mechanisms for our market, and analyse their theoretical properties. Then, in Section 4 we instantiate our model in two real-world scenarios. Using these scenarios, we evaluate and compare the proposed mechanisms empirically in Section 5. We conclude in Section 6.

2. AGENT MODELS

The system consists of a set of agents or *buyers*, $B = \{1, 2, \dots\}$, who arrive dynamically over time, and are interested in reserving a slot for charging their EV at one of the available charging stations, denoted by the set S . W.l.o.g., $i' > i$ means that buyer i' enters the market after i . Furthermore, we assume that charging occurs at discrete time slots (e.g., half-hourly slots), denoted by the set T , and that a car is fully charged during such a time slot (therefore, a buyer requires only a single time slot).⁵ Each buyer $i \in B$ can

⁵In future work, we plan to extend our model to include settings where buyers can partially charge and/or need several slots.

have different preferences regarding both the station he would like to charge at, as well as the time of the reservation. For example, in the park ‘n charge scenario, the buyer prefers a destination closer to his final destination. In the en-route charging scenario, the buyer prefers those stations which result in a smaller detour, and which are ideally placed between the departure point and the destination (e.g., if the battery’s state of charge is low, he would require a station close to the departure point). We abstractly represent such preferences by a matrix \mathbf{v}^i , where each element $v_{j,t}^i$ denotes the willingness to pay, or *value*, of an agent i for receiving a charging slot at time $t \in T$ in station $j \in S$. Note that this representation is very general, and can capture the costs (both in terms of time and money) due to a detour, as well as stations or times which are infeasible given the battery’s current state of charge (in which case the value for a particular slot or time is zero or even negative). These preferences constitute an agent’s private information (unknown to other buyers or the stations), also referred to as an agent’s *type*.

On the supply side, each station or *seller* can have multiple charging units, $K = \{1, 2, \dots\}$, which means that, for each particular time slot, possibly several reservations can be sold. Each station j has a cost for selling a certain number of time slots through the reservation system, denoted by the matrix \mathbf{c}^j , where $c_{t,k}^j$ is the cost of station $j \in S$ at time $t \in T$ for the k th unit, $k \in K$ (where $k = 1$ is the reservation for time slot t which has been allocated first, $k = 2$ the second reservation, etc.). If station j has at most k^j units available at a particular time t , we simply set $c_{t,k}^j = \infty$ for $k > k^j$. In practice, these costs represent an *opportunity cost*, i.e., the expected value of instead selling the unit on demand, without a reservation. This opportunity cost can be calculated by the probability that a certain unit is sold, multiplied by the profit of selling the unit on demand. Typically, peak times are expected to be more profitable (since the probability that the slot is used increases), and so have a higher opportunity cost. Furthermore, we assume that the marginal cost for additional units is non-decreasing. That is, $\forall j \in S, t \in T : c_{t,k+1}^j \geq c_{t,k}^j$. This is a natural assumption, as shown by the following example:

EXAMPLE 1. Consider a park ‘n charge with 1 time slot and 2 units. On-demand units are always sold at \$10. The station always manages to sell at least 1 unit on demand, and sells both units with probability 50%. Therefore, the opportunity cost of the first reservation is $0.5 \cdot \$10 = \5 , while the opportunity cost for selling the second unit through the reservation system is \$10.

These costs differ for each station, and are estimated by the stations based on observed past demand. Therefore, the costs constitute the station’s private information or *type*.

Given this, both buyers and sellers are asked to report their types to a *centre* (i.e., the reservation system) which then computes an allocation and payment for each agent. Buyers report their types as they enter the market, whereas sellers report their types in advance for the entire period T .⁶ Formally, let $\hat{\mathbf{v}}^i$ and $\hat{\mathbf{c}}^j$ denote the report for a buyer i and seller j respectively, $\hat{\mathbf{v}}$ and $\hat{\mathbf{c}}$ the reports of all buyers and sellers, and $\hat{\mathbf{v}}^{-i}$ ($\hat{\mathbf{c}}^{-j}$) all buyer (seller) reports except that of i (j). We define the *allocation* for buyer $i \in B$ by a tuple $x^i = \langle j, t, k \rangle$, if buyer i receives the k th unit, $k \in K$, of time slot $t \in T$ from seller $j \in S$ (note that k does not refer to a particular physical charging unit, but to the order in which the reservation was allocated), and use $x^i = \langle \emptyset \rangle$ to denote the case where the buyer is not allocated any slot. For the seller, we use \mathbf{x}^j to denote the

⁶In practice, the period can be limited to, e.g., the next 24 hours, and sellers can update their types as new time slots become available. For simplicity, we assume a single reporting stage.

allocation matrix at the end of the period T , where $x_{t,k}^j = 1$ means that the slot was allocated to a buyer, and $x_{t,k}^j = 0$ means the slot remained unallocated. Furthermore, we define p^i to be buyer i 's required payment to the centre, and p^j the payment received by seller j from the centre. We then define the utility function for a buyer $i \in B$ as:

$$U^B(\mathbf{v}^i, p^i, x^i) = \begin{cases} v_{j,t}^i - p^i & \text{if } x^i = \langle j, t, k \rangle \\ 0 & \text{if } x^i = \langle \emptyset \rangle \end{cases} \quad (1)$$

and for a seller $j \in S$ as:

$$U^S(\mathbf{c}^j, p^j, \mathbf{x}^j) = p^j - \sum_{(t,k) \in T \times K} c_{t,k}^j \cdot x_{t,k}^j \quad (2)$$

3. MARKET MECHANISM

Given the setting from the previous section, we would like to design a marketplace that satisfies the following properties:

- **Truthfulness:** This property requires that both sellers and buyers are incentivised to always truthfully report their types to the centre, i.e., reporting the true type is a dominant strategy.
- **No Delay:** In addition to truthful reporting, another way for buyers to strategise is to delay their entry into the market. No delay means that there is no incentive for buyers to do so.
- **Efficiency:** In general, an outcome is efficient if the goods are allocated to those who value them the most. Here, an efficient allocation maximises the sum of values minus the sum of costs.
- **Individual Rationality:** Individual rationality requires that participants are never worse off when participating in the mechanism. In this case, this means that both buyers and sellers have non-negative utilities.

- **Budget Balance:** *Weak* budget balance requires that the mechanism requires no outside subsidy, i.e., the sum of payments received from buyers is at least as great as the sum of payments to the sellers. The property is *strict* if they are equal.

In what follows, we first consider the allocation mechanism of the marketplace, which specifies how buyers are allocated to time slots and sellers as the buyers enter the marketplace. Then, we present various payment mechanisms for both buyers and sellers. Finally, we analyse the theoretical properties of both the allocation and payment mechanisms.

3.1 Allocation Mechanism

The allocation rule plays an important role in terms of obtaining efficiency (as well as the other properties). However, without perfect foresight, it is impossible to achieve perfect efficiency in online settings such as ours. Typically, researchers analyse worst-case results, and a simple greedy allocation rule has been shown to be 2-competitive (i.e., achieve at least half of the optimal) in settings with single-dimensional types [11, Ch. 16]. Furthermore, it has been shown to achieve close to optimal on average in related (albeit simpler) settings [6]. Given this, we also consider a greedy allocation mechanism which works as follows. Once the buyer arrives in the market and reports his type to the centre, he is immediately allocated the slot which maximises the difference between his value and the seller's cost, provided this is positive. Otherwise, no slot is allocated. More formally, an allocation x^i for buyer i is given by:⁷

$$x^i = f(\hat{\mathbf{v}}^i | \hat{\mathbf{c}}, X_i) = \arg \max_{(j,t,k) \in X_i | \hat{v}_{j,t}^i - \hat{c}_{t,k}^j \geq 0} \hat{v}_{j,t}^i - \hat{c}_{t,k}^j \quad (3)$$

⁷If there are multiple solutions, a tie-breaking rule is used.

where $x^i = \langle \emptyset \rangle$ if there is no solution and $X_i \subseteq S \times T \times K$ are the currently available allocations when buyer i enters the marketplace. Specifically, $X_1 = S \times T \times \{1\}$ for the first buyer entering the market and for subsequent buyers this is updated as follows:

$$X_{i+1} = \begin{cases} X_i \setminus \{\langle j, t, k \rangle\} \cup \{\langle j, t, k+1 \rangle\} & \text{if } x^i = \langle j, t, k \rangle \\ X_i & \text{if } x^i = \langle \emptyset \rangle \end{cases}$$

In addition, we will use $\mathbf{x}^j = g(\hat{\mathbf{c}}^j | \hat{\mathbf{c}}^{-j}, \hat{\mathbf{v}})$ to denote the seller allocation matrix as a function of her report, $\hat{\mathbf{c}}^j$.

3.2 Payment Mechanisms

Setting the payment correctly is important to obtain truthfulness, and to prevent speculation by strategic agents [11], which in turn should improve stability (reduce price fluctuations) and efficiency (since without truthful reports the allocations are likely to be inefficient). However, obtaining truthfulness on both sides of the market is a challenging problem and currently no truthful mechanism exists for the setting we consider. Furthermore, as we will show in Section 3.3, when using the greedy allocation mechanism discussed above, there exists no payment that always incentivises truthful reporting. Therefore, we only consider payments which incentivise truthful behaviour on the buyer side. On the seller side, there is no obvious choice of payment, and so we consider three different payments from the literature, each having different theoretical and empirical properties in our setting. In this section, we present these payments, while we discuss their theoretical properties in Section 3.3. We then go on to compare these payments experimentally in Section 5.

Buyer Payment: Whenever the buyer is allocated a slot, his payment is set equal to the seller's reported costs. Formally, $p^i = \hat{c}_{t,k}^j$ if $x^i = \langle j, t, k \rangle$, and $p^i = 0$ if $x^i = \langle \emptyset \rangle$.

For the seller payments, we consider three variants:

Posted Price: The seller payment for each allocated slot is equal to the reported cost for that slot. Therefore, the total payment is $p^j = \sum_{(t,k) \in T \times K} c_{t,k}^j \cdot x_{t,k}^j$. Note that this payment can be executed without the need for a centre, since sellers can simply post their prices for their currently available time slots, and each buyer agent then selects the best seller and time slot by solving Equation 3 and pays the posted price. Hence the name of the payment.

Reverse Vickrey: This seller payment is computed for each arriving buyer and is equal to the standard Vickrey payment for a single item, albeit for a reverse setting (with a single buyer and multiple sellers). Formally, let $X_i^{-j} = X_i \cap S \setminus \{j\} \times T \times K$ denote the available allocations when buyer i arrives in the market, but without the slots from seller j . Then, let $x^i = \langle j, t, k \rangle = f(\hat{\mathbf{v}}^i | \hat{\mathbf{c}}, X_i)$ denote the actual allocation and $x_{-j}^i = \langle j', t', k' \rangle = f(\hat{\mathbf{v}}^i | \hat{\mathbf{c}}^{-j}, X_i^{-j})$ the best alternative allocation without considering seller j . Given this, seller j 's payment for an allocation x^i is:

$$p_{t,k}^j = \begin{cases} 0 & \text{if } x^i = \langle \emptyset \rangle \\ v_{t,k}^i & \text{if } x^i \neq \langle \emptyset \rangle \wedge x_{-j}^i = \langle \emptyset \rangle \\ v_{t,k}^i - (v_{t',k'}^i - c_{t',k'}^j) & \text{otherwise} \end{cases} \quad (4)$$

and the total payment is: $p^j = \sum_{(t,k) \in T \times K} p_{t,k}^j \cdot x_{t,k}^j$.

Critical Value: The notion of critical value has been introduced in the mechanism design literature to produce truthful mechanisms for *single-parameter domains*, in which the allocation decisions are binary (the agent is either allocated the item or bundle, or not) and an agent simply has a value for 'winning' (see, e.g., [11, Ch.9]). In these domains, the critical value is equal to the lowest value (or highest cost) that could have been reported and still wins the item or bundle (i.e., for which the allocation remains unchanged). In our

setting, however, we need to adapt the notion of critical value since a seller has possibly multiple slots and multiple units for each slot, and so her type is multi-dimensional.⁸ Specifically, given an allocation \mathbf{x}^j , we define the critical value in our setting as the highest *sum of costs* for the allocated slots, for which the allocation remains unchanged in terms of the slots sold. That is, the slots and units can be sold to different buyers, as long as the allocation matrix, \mathbf{x}^j , remains the same. Formally:

$$p^j = \max_{\hat{\mathbf{c}}^j \in C^j | g(\hat{\mathbf{c}}^j | \hat{\mathbf{c}}^{-j}, \hat{\mathbf{v}}) = \mathbf{x}^j} \sum_{(t,k) \in T \times K} c_{t,k}^j \cdot x_{t,k}^j \quad (5)$$

where C^j is the set of all possible reports of type from seller j . Note that, to compute the payment, the centre needs to rerun the entire market for each possible seller misreport. Furthermore, in the case each seller only has a single time slot and unit, the above equation reduces to the critical value for single-valued domains as in [11, Ch.9] (albeit for *costs* instead of *values*).

3.3 Theoretical Properties

We now discuss the theoretical properties of the allocation and payment mechanisms discussed in Section 2. We first consider the properties for a buyer, then for a seller and, finally, for the centre. For the buyer, three properties apply: truthfulness, no delay, and individual rationality. It is easy to see that, given the right conditions and using the buyer payment, all these properties are satisfied:

THEOREM 1. *The Buyer Payment as defined in Section 3.2 is truthful and individually rational for buyers. Moreover, if $\forall j \in S, t \in T, k \in K : c_{t,k+1}^j \geq c_{t,k}^j$, buyers have no incentive to delay their entry into the market, i.e., no delay.*

PROOF. It is evident that Buyer Payment is individually rational as $p^i \leq v_{j,t}^i$ if $x^i = \langle j, t, k \rangle$ according to Equation 3 and therefore buyer i 's utility $U^B(\mathbf{v}^i, p^i, x^i) \geq 0$ (see Equation 1).

According to Proposition 9.27 of [11], the payment is truthful for buyers because (i) p^i does not depend on \mathbf{v}^i and given other agents' reports, p^i is the same for any \mathbf{v}^i that gives the same allocation x^i , (ii) the greedy mechanism (i.e., Equation 3) optimises for each buyer i . Since $\forall j \in S, t \in T, k \in K : c_{t,k+1}^j \geq c_{t,k}^j$, so $\max_{(j,t,k) \in X_i} v_{j,t}^i - c_{t,k}^j$ might decrease if i enters later, given that i 's valuation remains the same or decreases when i enters later. Thus, $U^B(\mathbf{v}^i, p^i, x^i)$ might decrease when i delays his entry. \square

In terms of the sellers, all the payments proposed satisfy individual rationality (we omit a formal proof since the results are trivial). However, as mentioned in Section 3.2, obtaining truthfulness is challenging and none of the proposed payments satisfy this property in general. Indeed, as the following theorem shows, when using the greedy allocation mechanism from Section 3.1, there *exists no payment* that always satisfies truthfulness.

THEOREM 2. *If slots are allocated using the greedy mechanism given by Equation 3, there exists no payment which always satisfies truthfulness for the seller.*

PROOF. The proof is based on showing, by example, that the allocation rule violates weak monotonicity (WMON), which is a necessary condition for such a payment to exist (see [2] for details). Let $V(\mathbf{c}^j, \mathbf{x}^j) = -\sum_{(t,k) \in T \times K} c_{t,k}^j \cdot x_{t,k}^j$ denote seller j 's utility/valuation without considering any payments (noting that this

⁸Note that [2] defines payments for multi-dimensional types. However, these payments are only applicable to (partially) ordered domains, which is not the case in our setting. We refer the reader to [2] for further details.

is always negative or zero). Given this, WMON requires that, for all $\mathbf{c}^j, \hat{\mathbf{c}}^j \in C^j$, the following holds (we remove the conditional part of the seller allocation function, g , for brevity):

$$V(\hat{\mathbf{c}}^j, g(\hat{\mathbf{c}}^j)) - V(\mathbf{c}^j, g(\hat{\mathbf{c}}^j)) \geq V(\hat{\mathbf{c}}^j, g(\mathbf{c}^j)) - V(\mathbf{c}^j, g(\mathbf{c}^j)) \quad (6)$$

Now, consider an example with 1 seller, 2 time slots, and 1 unit per time slot. The seller's costs are given by $\mathbf{c}^1 = \langle 2, 2 \rangle$ (i.e., costs are equal for both time slots). Furthermore, there are two buyers, the first one entering the market has type $\mathbf{v}^1 = \langle 6, 5 \rangle$, while the second one has type $\mathbf{v}^2 = \langle 10, 0 \rangle$. Given truthful reporting, the greedy mechanism will allocate the first buyer to the first slot, and not allocate the second buyer (since the early slot is taken). Therefore, $V(\mathbf{c}^1, g(\mathbf{c}^1)) = -2$. Now, if the seller's costs for the two slots were to be *higher*, e.g., $\hat{\mathbf{c}}^1 = \langle 8, 3 \rangle$, she would be allocated both slots, and $V(\hat{\mathbf{c}}^1, g(\hat{\mathbf{c}}^1)) = -11$. Furthermore, we have that $V(\mathbf{c}^1, g(\hat{\mathbf{c}}^1)) = -4$ and $V(\hat{\mathbf{c}}^1, g(\mathbf{c}^1)) = -8$. Given these values, Equation 6 becomes: $-11 + 4 \geq -8 + 2 \leftrightarrow -7 \geq -6$, which clearly does not hold. \square

The intuition for the above impossibility theorem is as follows. Since the centre has no perfect foresight and greedy allocation is necessarily inefficient, it sometimes allocates *more* buyers to a seller if she reports *higher* costs, thereby violating WMON. An anticipating seller could always take advantage of such a situation, and so the mechanism is not truthful in dominant strategies. This problem can be addressed by "ironing" the allocation, which involves cancelling those allocations which violate WMON. However, even for relatively simple settings, such as single-parameter domains [12] or one-sided markets where buyers have decreasing marginal values [6], these cancellations introduce additional inefficiencies (since the canceled units cannot be reallocated), are generally computationally demanding, and are often difficult to apply in practice (e.g., in the EV charging scenario, it requires discharging a car). Instead, we take a different approach and choose to relax the truthfulness property. Specifically, we consider weaker notions of truthfulness in the hope that these will reduce (but not eliminate) speculation in practice. In Section 5 we then empirically compare the mechanisms by allowing agents to strategise.

In particular, the following theorem shows that the Reverse Vickrey payment (see Section 3.2) satisfies *myopic truthfulness* for sellers, which we define as being truthful for a seller when ignoring future buyers. Formally:

THEOREM 3. *Suppose that a seller j is able to report a cost matrix, \mathbf{c}^j , for each buyer i entering the market. Then, Reverse Vickrey is truthful if i is the last buyer and this is known by j .*

PROOF. The assumption in the theorem leads to an equivalent case where there is only one buyer buying one time slot from multiple sellers. It is evident that (i) p^j does not depend on \mathbf{c}^j , and it is the same for any \mathbf{c}^j that receives the same allocation for j given other agents' reports, and (ii) j 's utility $U^S(\mathbf{c}^j, p^j, \mathbf{x}^j)$ is maximised because of Equations 3 and 4. Therefore, sellers are truthful under this assumption. \square

Furthermore, we show that the Critical Value payment is truthful for a specific setting where each buyer is only interested in a single (but possibly different) time slot — buyers are said to be *single-minded* in terms of time slots — and each seller has at most a single unit for each time slot (in Section 4.1 we discuss these assumptions in relation to the park 'n charge scenario). Formally:

THEOREM 4. *If $\|K\| = 1$ and $\forall i \in B, \forall j, j' \in S, \forall t, t' \in T : v_{j,t}^i > 0 \wedge v_{j',t'}^i > 0 \Rightarrow t = t'$, Critical Value is truthful.*

PROOF. (Sketch) Since each buyer is interested in at most one time slot, the seller cannot influence to which slot a buyer is allocated by misreporting. Therefore, in terms of incentives, the setting is equivalent to one where each time slot is sold by independent sellers. Given that there is at most 1 unit, each such seller’s type can be represented by a single parameter. It is easy to see that, in this setting, greedy allocations are monotonic, and so the critical value payment incentivises truthful bidding in single-parameter domains (see also [11, Ch. 9.5.4]). \square

The following example shows that truthfulness cannot be achieved for more general settings using the Critical Value payment.

EXAMPLE 2. Consider a seller with 1 time slot and 2 units, and type $\mathbf{c}^1 = \langle 1, 5 \rangle$ for her two units, and two buyers with $\mathbf{v}^1 = \langle 4 \rangle$ and $\mathbf{v}^2 = \langle 6 \rangle$. If the seller reports her true type, both units are allocated and her critical payment is 10 (as she could have reported up to $\langle 4, 6 \rangle$), and her utility is $U^S = 10 - 6 = 4$. On the other hand, if she reports $\mathbf{c}^1 = \langle 5, 5 \rangle$, only the first slot is allocated, and her critical payment is 6. However, her utility is $6 - 1 = 5$ and so she is better off misreporting.

Similarly, we can use the example in the proof of Theorem 2 to show that truthfulness is violated when bidders are not single-minded in term of time slots. In this example, the critical payment is 10 when the seller is truthful (the value of buyer 2), and her utility is $U^S = 10 - 2 = 8$. On the other hand, if the seller reports between 6 and 10 for the first slot, both slots get allocated, the critical payment is 15, and her utility is $U^S = 15 - 4 = 11$. Clearly, the seller has an incentive to misreport.

So far, we have not discussed the Posted Price payment. Trivially, this payment is not truthful for any of the above special cases. However, of the three seller payments, only Posted Price is budget balanced (since, for each allocation, the payment received from the buyer is equal to the payment given to the seller), which is an advantage for the centre. For the other two payments, the centre always runs a *deficit* (since the payments to the sellers are always at least as high as the payments received from the buyers).⁹

4. REAL-WORLD SCENARIOS

The model presented so far contains an abstract representation of buyer valuations. To demonstrate how buyer valuations can be derived in practice, and the role of the buyer agent in doing so, we present instantiations of two different scenarios: park ’n charge (where the main aim is to find the charging station closest to the final destination), and en-route charging (where the main aim is to arrive at some destination at a particular point in time). These scenarios are then used in Section 5 to evaluate the payment schemes proposed within a realistic setting.

4.1 Park ’n Charge Scenario

In this scenario, we consider a situation where EVs charge while parked away from home for an extended period of time (e.g., a full day), and where users can reserve a parking place (i.e., charging station) ahead of time. This can include large car parks near shopping centres as well as smaller, private car parks.¹⁰ Since the

⁹In practice, the budget balance problem can be addressed by charging fees, such as participation fees for buyers and/or sellers. However, fees are outside the scope of this paper. Nevertheless, in Section 5 we experimentally compare the size of the deficit for different payment mechanisms and scenarios.

¹⁰Several services for reserving a parking space already exist in practice, such as www.parkatmyhouse.com, although they currently do not include charging facilities.




Figure 1: Example for en-route charging scenario.

parking is for an extended period, the buyer is typically interested in a particular time slot (e.g., a certain day), and so it is reasonable to assume that a buyer is single-minded in terms of the time slots. At the same time, a buyer has different preferences about the locations of the parking place. Specifically, we assume that a buyer’s utility depends on the travel distance to the station, as well as the convenience in terms of reaching the actual destination. In the experiments, we use the following simple linear value function to capture these aspects:

$$V_i(d^{\text{drive}}, d^{\text{walk}}) = v_i^{\text{max}} - \alpha_i \cdot d^{\text{drive}} - \beta_i \cdot d^{\text{walk}} \quad (7)$$

where d^{drive} represents the distance from the point of origin to the charging station, d^{walk} the distance from the parking place to the actual destination (which can be walking distance or, depending on the exact setting, the distance or time using local public transport), and α_i and β_i capture the user’s preferences in terms of two distances. Furthermore, v_i^{max} represents the cost of the buyer’s most preferred *outside option* when not using the reservation system (e.g., the cost of an on-demand parking place or using an alternative mode of transportation). This cost can include the actual price paid for using the outside option, as well as inconveniences such as increased travel time. This cost is avoided if the agent is allocated a slot, and so it represents a *value* in Equation 7.

Now, using this function, the user only requires to specify the preference parameters α_i , β_i , and v_i^{max} whereas the agent is able to compute the d_j^{drive} and d_j^{walk} for each station j by using a routing algorithm and other available information (such as travel time using local transport). Given this and the required time slot, t , the agent can compute the buyer value matrix by setting $v_{j,t}^i = V_i(d_j^{\text{drive}}, d_j^{\text{walk}})$, $j \in S$ and $v_{j,t'}^i = 0$ for $t' \neq t$.

4.2 En-Route Charging Scenario

In this scenario, we assume that a buyer has a preferred time of arrival at the final destination, denoted by t_i^* (e.g., arriving at work at 9:00am), at which his utility is maximised. The utility decreases if the true arrival time is either too early or too late. Furthermore, the buyer prefers shorter car journeys (including the charging) to longer ones. For illustration purposes, we model these preferences by simple constants, which indicate a linear decay in utility, although other functions can be easily applied. Formally, the value derived given a departure time, t^{dep} , and actual arrival time at destination, t^{arr} , is given by:

$$V_i(t^{\text{dep}}, t^{\text{arr}}) = v_i^{\text{max}} - \alpha_i \cdot (t^{\text{arr}} - t^{\text{dep}}) - c(t_i^* - t^{\text{arr}}) \quad (8)$$

where $c(x) = \beta_i \cdot x$ if $x \geq 0$, or $c(x) = -\gamma_i \cdot x$ if $x < 0$. Here, α_i is a coefficient modeling the buyer’s cost for the journey time, while $\beta_i, \gamma_i \geq 0$ model the costs for early and late arrival, respectively. As before, v_i^{max} represents the *cost* of the buyer’s most preferred outside option when not using the reservation system.

Now, given an appropriate routing algorithm, the agent can compute the required departure time, t^{dep} , and arrival time at destination, t^{arr} , when using a particular station j and time slot t , which can then be used to compute the value of that station and time slot, $v_{j,t}^i$. Consider the simple example in Figure 1, which contains 2

stations, s_1 and s_2 with the same 2 time slots each, t_1 and t_2 . The values on the edges represent the travel time in minutes. If we assume that charging takes 30 minutes, the values can be computed as follows: $v_{s_1,t_1}^i = V_i(t_1 - 45, t_1 + 60 + 30)$, $v_{s_1,t_2}^i = V_i(t_2 - 45, t_2 + 60 + 30)$, $v_{s_2,t_1}^i = V_i(t_1 - 75, t_1 + 20 + 30)$, and $v_{s_2,t_2}^i = V_i(t_2 - 75, t_2 + 20 + 30)$.

Note that Equation 7 could be easily extended to take into account other factors, such as deadlines and earliest departure times, which would exclude certain time slots from being feasible. In addition, in practice each vehicle battery has a certain *state of charge* (SOC), which limits the driving range and so the feasible stations. Similar to the driving distance, the agent can use a navigation system to calculate the required SOC for reaching a particular station (most existing satellite navigation systems already have this ability which is used to allow users to select the most fuel efficient route). Finally, we note that both Equations 7 and 8 are illustrative examples and could be replaced by more complex, non-additive utility functions without affecting the mechanisms discussed in this paper.

5. EVALUATION

In this section, we evaluate our proposed payment mechanisms empirically, in order to examine and critically compare them in realistic settings. Specifically, while we discussed several theoretical properties in Section 3.3 (including whether they are budget balanced, truthful and efficient), we are now interested in quantifying some of their performance characteristics in practice (such as how large their deficits are, how strategic behaviour affects the system and how far they are from optimal efficiency). In the following, we first describe our experimental methodology, and then give results for the two real-world scenarios we described in Section 4.

5.1 Experimental Methodology

To evaluate our payment mechanisms in practice, we simulate a number of realistic settings by randomly sampling buyers and sellers from certain distributions (detailed below). When the mechanism is truthful, as is the case for buyers and for Critical Value payments to stations in the settings defined by Theorem 4, we assume that participants adopt the dominant strategy and report truthfully.

In all other cases, sellers may benefit from strategically misreporting their costs. In order to capture this strategic behaviour in a principled manner, we use a simple form of adaptive learning, iterated best response [10]. Using this learning approach, we first sample the set of buyers and sellers, and then initialise the reports, \hat{c}^j , of each seller j to random values. To ensure computational feasibility, we restrict the possible reports for each cost $\hat{c}_{t,k}^j$ to a finite set,¹¹ and we enforce $\hat{c}_{t,k}^j \leq \hat{c}_{t,k+1}^j$. Then, we iteratively choose each seller in sequence and set her report to the best response, given the current reports of others (if this improves on her previous action). When a seller is indifferent between several better responses, we choose the one with the highest sum of costs for allocated slots, as we found this speeds up convergence in practice. This continues until no seller has an incentive to change her response, at which point the reports correspond to a Nash equilibrium [10]. We also terminate the search when it has iterated over all sellers 100 times, which typically indicates that they have entered a cycle of best responses that do not converge.

We chose iterated best response here, because it is one of the simplest forms of learning, where agents act only in response to the utilities they derive and without anticipating the future responses of others. Thus, if this approach converges quickly to a Nash equi-

¹¹In the experiments, we choose 11 fixed costs that evenly discretise the distribution of buyer valuations.

librium, it is an indication that agents do not require sophisticated learning techniques to perform well in these settings and that stable equilibrium prices can be reached through simple reactive strategies. Note, however, that this analysis assumes agents can compute their best responses iteratively for a given, fixed set of buyers. In practice, buyer populations change from day to day, and stations have limited knowledge about the consequences of their actions, which they may need to explore over time. In this paper, we abstract away from these complexities and focus on a principled equilibrium analysis, leaving other learning approaches to future work.

Given this methodology and in order to compare our alternative payment mechanisms, we record the following system properties when equilibrium has been reached (or after 100 iterations):

- **Efficiency:** The difference between all buyer values and station costs, as a proportion of the optimal. Specifically, this is $W(\mathbf{x})/W(\mathbf{x}^*)$, where \mathbf{x} is the vector of all buyer allocations chosen by the mechanism, $W(\mathbf{x}) = \sum_{i \in B} w(i, \mathbf{x}^i)$ and $w(i, \langle j, t, k \rangle) = v_{j,t}^i - c_{t,k}^j$ (with special case $w(i, \langle \emptyset \rangle) = 0$). The vector \mathbf{x}^* is the optimal allocation that maximises $W(\mathbf{x})$.

- **Convergence:** Whether the iterated best response algorithm converged to an equilibrium within 100 iterations. As discussed above, this is an indicator for the stability of the mechanism.

- **Deficit:** The overall difference between the payments received from buyers, and the payments made to sellers, as a proportion of the optimal values derived, i.e., $(\sum_{i \in B} p_i - \sum_{j \in S} p_j)/W(\mathbf{x}^*)$.

For statistical significance, we repeat all experiments 1000 times and show 95% confidence intervals. We also consider settings with varying numbers of buyers, keeping the number of stations fixed, in order to investigate different ratios between supply and demand.

5.2 Park 'n Charge Results

To simulate the park 'n charge scenario outlined in Section 4.1, we assume that stations and buyers are distributed on a two-dimensional Euclidean plane, and that the travel distance between two points is equal to the straight-line distance between them (where the unit length is 1km). More specifically, we assume that some area of interest is centred on coordinates $(0, 0)$, e.g., the centre of a city, a shopping district or a business park. We then generate the locations of 15 charging stations around this by adding Gaussian noise with standard deviation $\sigma = 1$ km independently to the x and y coordinates of the central location. This approximately follows the patterns observed on PlugShare, where home charging locations are clustered around urban centres. We assume costs of all stations are 0 and, for simplicity, we assume $\|T\| = 1$ and $\|K\| = 1$. This mirrors a setting where stations are private home owners that rent out a single parking space, and it ensures the Critical Payment mechanism is truthful for sellers (see Section 3.3).

To simulate buyers, we assume they start their journey further away from the centre of interest (we add noise with $\sigma = 50$ km), while their destination is closer to the centre (perturbed with $\sigma = 1$ km). Thus, their journey represents a considerable drain on the EV battery. To instantiate the values for v_i^{\max} , α_i and β_i in Equation 7 for a particular buyer i in an intuitive manner, we first assume an ideal reference station that is located exactly on the buyer's destination coordinates. We then generate a (hypothetical) value v_i^{ref} for this station from the uniform distribution shown in Table 1(a). Next, we randomly determine the ratio of the two penalty parameters β_i/α_i . This indicates the buyer's relative preference for driving over walking, e.g., $\beta_i/\alpha_i = 10$ indicates the buyer is indifferent between driving an additional 10km or walking an additional 1km. Finally, we determine a maximum walking distance d_i^{\max} the buyer is willing to tolerate (assuming d^{drive} is equal to the direct

(a) Park 'n charge.

Variable	Distribution
v_i^{ref}	$\mathcal{U}(\$5, \$20)$
β_i/α_i	$\mathcal{U}(2, 15)$
d_i^{max}	$\mathcal{U}(0.5, 5)\text{km}$

(b) En-Route Charging.

Variable	Distribution
v_i^{ref}	$\mathcal{U}(\$0, \$10)$
t_i^*	$\mathcal{U}(30, 120)\text{min}$
t_i^{earliest}	$t_i^* - \mathcal{U}(60, 180)\text{min}$
t_i^{latest}	$t_i^* - \mathcal{U}(30, 120)\text{min}$
λ_i	$\mathcal{U}(0.5, 2)$

Table 1: Experimental variables.

line distance to the destination) before preferring the outside option. These parameters together allow us to calculate v_i^{max} , α_i and β_i , and therefore V_i , for any particular charging station. When the buyer has a negative value for all stations, we re-sample that buyer.

Figure 2 shows the full results in the park 'n charge scenario as we vary the number of buyers. Here, we compare the outcomes of our different mechanisms when all stations report truthfully (labeled *truthful* in the figure), and when they strategise using iterated best response (labeled *learning*). The first significant result here is that the efficiency of the mechanisms is far worse when sellers report their costs truthfully, but reaches 90% or more of optimal when stations strategise. This is because the greedy allocation rule will allocate all buyers in sequence (where feasible), regardless of whether buyers with higher valuations may arrive in the future. When stations strategise in the Posted Price and Reverse Vickrey mechanisms, however, they artificially increase their reported costs to ensure they are allocated high-value agents (and thus receive higher payments).

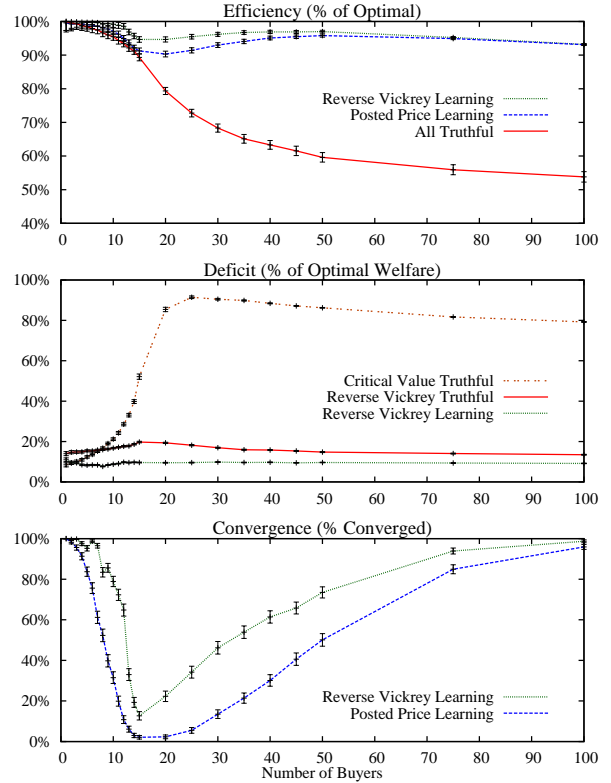
Interestingly, the truthful Critical Value mechanism does not perform very well. In addition to receiving a low efficiency when there are many buyers in the market, it incurs a high overall deficit. This is because the mechanism effectively compensates stations for being allocated low-value agents, which is in some cases more than the overall value generated. The Reverse Vickrey mechanism also incurs a deficit, but this is relatively low. Specifically, in the strategic setting, this is around 10% and lower than the same mechanism in the truthful setting, because stations generally report higher costs, leading to a higher revenue generated from buyer payments.

Finally, we note that iterated best response converges consistently to a Nash equilibrium in the Posted Price and Reverse Vickrey mechanisms when demand is low or very high. Between these extremes, convergence drops, with a minimum approximately when the number of buyers equals the number of stations. Briefly, this is because the valuations of buyers are sufficiently diverse here to allow stations to repeatedly undercut their competitors' prices until some stations are forced out of the allocation, prices are raised again and the cycle repeats. However, the Reverse Vickrey mechanism often leads to more stable outcomes. This is because the mechanism has less scope for strategic behaviour — a station's reported cost influences what buyer she is allocated, but not the payment she receives through that allocation. The better convergence properties of the Reverse Vickrey mechanism also explain the small increase in efficiency over the Posted Price mechanism. When a setting fails to converge, it typically results in a poorer allocation.

5.3 En-Route Charging Results

Next, we look at the en-route charging scenario. To ensure that calculating the best response is computationally feasible here, we restrict our analysis to 3 stations, with 2 possible time slots (each of which corresponds to 30 minutes of charging) and 2 charging units.¹² This time, we assume buyers are wishing to travel between two areas on a Euclidean space, one of which is centred

¹²In practice, agents can employ heuristics to deal with far larger settings, and we have verified that similar trends are observed when

**Figure 2: Results of the park 'n charge scenario.**

on $(-100, 0)$ and the other one on $(100, 0)$ (i.e., 200km apart), with Gaussian noise ($\sigma = 10\text{km}$) added independently to each coordinate. This represents a long-distance journey that will likely exhaust the EV battery, necessitating a stop at a recharging station. Stations are again centred on $(0, 0)$, with noise added ($\sigma = 50\text{km}$). As stations in these settings may be able to sell their charging capacity to customers without reservations, we assume stations occur opportunity costs, which we draw independently from the uniform distribution $\mathcal{U}(0, v^{\text{high}})$, where v^{high} is the expected maximum valuation of any customer.

To generate v_i^{max} , α_i , β_i and γ_i , we again assume an ideal station that immediately provides the buyer with a full charge at his departure location, and we attach a random value v_i^{ref} to this (see Table 1(b) for the distributions). We also draw a random deadline t_i^* , which we represent in minutes after the start of the first charging slot. Next, we draw random parameters for the earliest time the buyer is willing to arrive at the destination, t_i^{earliest} , as well as the latest acceptable arrival time, t_i^{latest} . Specifically, assuming a direct drive from departure point to destination, these are the times at which the buyer becomes indifferent to taking the outside option. To capture the driver's willingness to take a detour, we sample a delay tolerance, λ_i , which is the maximum proportional increase over the direct travel time that the driver is willing to tolerate (assuming arrival at t_i^*). For example, when $\lambda_i = 0.75$ and the usual direct travel time is 80 minutes, the driver becomes indifferent to the outside option when the detour to a charging station (including charging) takes an additional hour. Finally, we sample the remaining state of charge uniformly at random, such that the vehicle can reach at least one charging station, but not the destination. Assuming an average travel speed of 100 km/h, these parameters are

using simple hill-climbing to approximate the best response.




Figure 3: Results of the en-route charging scenario.

again sufficient to calculate v_i^{\max} , α_i , β_i and γ_i for buyer i , and we re-sample any buyers that do not have any $v_{j,t}^i > 0$ (or where the destination is closer to the departure location than the nearest charging station).

Figure 3 contains the results for the en-route charging setting. The overall trends are broadly similar here to the park 'n charge scenario, so we concentrate on the main differences. First, we additionally include the Critical Value payment mechanism with learning agents here, because it is no longer truthful for stations. Perhaps surprisingly, the mechanism performs very poorly here, compared to the other non-truthful mechanisms. In particular, efficiency suffers, because stations no longer have an incentive to be allocated late-arriving high-value buyers, as the mechanism will later correct for these in any case. As a consequence of this, the deficit also rises when there are many buyers — stations can report low costs, resulting in low payments from the buyers, but still collect high critical value payments from the mechanism.

Next, we note that overall convergence for the Posted Price and Reverse Vickrey mechanisms is significantly higher than in the park 'n charge scenario. This is due to the introduction of positive costs in this setting, which significantly reduce the scope for strategising, as stations may incur losses when they under-report their costs.

In conclusion, the results throughout this section indicate that strategic behaviour on the seller side typically leads to a significantly higher overall efficiency than truthful mechanisms (where these are feasible, as in the park 'n charge scenario, or where agents do not strategise). Furthermore, the results show that a simple Posted Price mechanism leads to a high efficiency and may be the preferred choice because it is budget balanced. In some settings, where stability of the prices is important, the Reverse Vickrey is advantageous, and can even result in a small increase in efficiency, but it runs at a slight deficit.

6. CONCLUSIONS

In this paper, we considered the problem of allocating slots at charging stations to EV owners. To address this, we proposed a novel advance reservation system that matches drivers to available stations using a two-sided market, and we explored a number of potential payment mechanisms for both sides of this market. Using a principled equilibrium analysis, we showed that strategic behaviour on the seller side can lead to a high overall efficiency, and we demonstrated that some of our mechanisms offer a trade-off between stability and budget balance. In practice, we envisage that our mechanism can be integrated with assistive driver technologies that seamlessly perform both routing and instant advance reservation functions on behalf of the driver.

In future work, we plan to explore some of the human challenges associated with participating in the market we propose, and, in particular, how some of its complexities can be hidden from users. We will also explore more complex scenarios, including long-distance trips with multiple recharging stops and partial charging, and we will investigate other learning approaches beyond iterated best response that can deal with uncertainty in demand.

Acknowledgement: This work was supported by the ORCHID (orchid.ac.uk) and iDEaS projects (ideasproject.info).

7. REFERENCES

- [1] *Electric Vehicles: Charged with potential*. Royal Academy of Engineering, 2010.
- [2] 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.
- [3] A. Blum, T. Sandholm, and M. Zinkevich. Online algorithms for market clearing. *J. ACM*, 53(5):845–879, Sept. 2006.
- [4] J. Bredin, D. C. Parkes, and Q. Duong. Chain: A Dynamic Double Auction Framework for Matching Patient Agents. *J. of Artif. Intel. Research*, 30:133–179, 2007.
- [5] 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.
- [6] E. H. Gerding, V. Robu, S. Stein, D. C. Parkes, A. Rogers, and N. R. Jennings. Online mechanism design for electric vehicle charging. In *AAMAS'11*, 2011.
- [7] M. Gonen, R. Gonen, and E. Pavlov. Generalized trade reduction mechanisms. In *ACM EC'07*, pages 20–29, 2007.
- [8] H.-J. Kim, J. Lee, and G.-L. Park. Constraint-based charging scheduler design for electric vehicles. In *Intelligent Information and Database Systems*, Springer LNCS vol. 7189, pages 266–275, 2012.
- [9] R. P. McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56(2):434–450, 1992.
- [10] P. Milgrom and J. Roberts. Adaptive and sophisticated learning in normal form games. *Games and Econ. Behavior*, 3:82–100, 1991.
- [11] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. *Algorithmic Game Theory*. Cambridge Univ. Press, 2007.
- [12] D. C. Parkes and Q. Duong. An ironing-based approach to adaptive online mechanism design in single-valued domains. In *AAAI'07*, pages 94–101, 2007.
- [13] H. Qin and W. Zhang. Charging scheduling with minimal waiting in a network of electric vehicles and charging stations. In *Proc. 8th ACM Int. Workshop on Vehicular Inter-Networking*, pages 51–60, 2011.
- [14] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings. Putting the 'smarts' into the smart grid: a grand challenge for artificial intelligence. *Commun. ACM*, 55(4):86–97, 2012.
- [15] S. Stein, E. Gerding, V. Robu, and N. R. Jennings. A model-based online mechanism with pre-commitment and its application to electric vehicle charging. In *AAMAS'12*, pages 669–676, 2012.
- [16] M. Vasirani and S. Ossowski. A computational monetary market for plug-in electric vehicle charging. In *AMMA'11*, 2011.