FleetPower: Creating Virtual Power Plants in Sustainable Smart Electricity Markets

Kahlen, Micha
Rotterdam School of Management, Erasmus University
kahlen@rsm.nl

Ketter, Wolfgang
Rotterdam School of Management, Erasmus University &
University of Cologne
wketter@rsm.nl

Gupta, Alok
Carlson School of Management, University of Minnesota
gupta037@umn.edu

Introduction

Carsharing is increasingly becoming a popular business model (Firnkorn and Müller, 2011) where cars are rented for short durations (sometimes minutes). Unlike traditional car rentals, where renters keep possession of the car even during the long unused hours, the idea of carsharing is to rent the cars for short, one way trips. Many of the carsharing companies use electric vehicles (EV) as their carsharing fleets. Clearly, while the cars are parked the carsharing companies incur investment and maintenance costs. While this loss is unavoidable for traditional fossil fuel based cars, for fleets with EVs additional revenues can be generated by appropriate use of EVs as virtual power plant (VPP), a collection of distributed power sources that are centrally coordinated by an information system (IS) to offset energy imbalances (Pudjianto et al., 2007). In this paper, we develop and validate a method to increase the utilization and hence the profits of electric vehicle (EV) carsharing fleets. Our approach allows EVs to charge when there is a surplus of electricity and discharge it to the grid (vehicle-2-grid, V2G) when there is a shortage of electricity. We develop a computational control mechanism for the VPP to decide how the EVs should be allocated over time in terms of charging, discharging, and being available for rental customers. Based on a discrete event simulation model that is calibrated using real data on availability of vehicles and their movements in three distinct locations supplemented with data on electricity prices on reserve markets, we show what extra profits can be generated from EVs in addition to the rental business, by participating in these electricity auction markets. In essence, we develop *smart markets* (McCabe et al., 1991; Gallien and Wein, 2005; Bichler et al., 2010) to enable energy sources to be deployed cost-effectively. Fleets can offer the storage on short term electricity markets. They can chose between the day-ahead, intraday, and operating reserve market based on the largest price differentials. In this research we seek to analyze the operation in control reserve market, which acquires back-up power sources that can consume electricity from or dispatch electricity to the grid within seconds. These back-up sources guarantee an alternative source of power when another power source produces more or less then had been promised (for example, due to technical defects or weather-related issues). The fact that EV batteries can be charged, and discharged flexibly makes them very suitable for offering control reserve power at economic rates. However, the findings in this paper are therefore a conservative estimate of actual profits, as we only look at the price differentials in one markets.

To operationalize our model, we compute trading prices (bids and asks) for every 15-minute time interval of a given week to charge and discharge EVs and these prices are offset against the opportunity costs of not renting out the EV. We have develop an intelligent software agent (Collins et al., 2010), called *FleetPower* which combines information from the location, the state of charge of the battery (SoC), historical rental transaction data, and historical prices on the control reserve market and uses this information to make optimal allocation of EVs to either the rental market or a VPP (charging or discharging). The agent considers both the profits from charging and discharging and what effect the withdrawal of vehicles for rent might have on the mobility of those wanting to hire cars (sociotechnical implications). This means that our model makes an explicit trade-off between the asymmetric benefits to be gained from either offering cars for rental or using them to balance the grid in real time. To generate profits for the fleet, the system optimizes the allocation of EVs either to those needing to rent an EV (the social part) or to the balancing market which purchases services to balance the grid (the technical part). The focus is on ensuring that the availability of EVs for rental is not compromised, as the opportunity cost of loosing a rental customer is very high. While the profits from renting out a vehicle are on average around $15 per transaction, profits from the charging and discharging are only a few cents per 15 minutes. As a consequence FleetPower offers the VPP capabilities conservatively.

We use real data from EV carsharing fleets and control reserve markets to test our strategies via a simulation platform that allows us to calibrate the trading prices. The actual data on electric vehicle carsharing fleets was provided by Daimler's subsidiary, Car2Go. Car2Go is a city carsharing service where customers can rent cars, which they pay for on a per-minute basis. To use the service users have to register and pay a one-time registration fee. Once registered, they can make use of the freefloat service which enables them to pick up and drop off vehicles anywhere within the city boundaries, not necessarily in the same location. The service is available in 30 major cities in North America and Europe, with a total of more than 13,000 Smarts ForTwo. We tracked the location, state of charge, and transactions of 500 EVs in Stuttgart, 300 in Amsterdam, and 300 in San Diego for 14 months. These cities have 220, 1,500, and 100 charging stations respectively. The high number of charging stations in Amsterdam is related to government incentives. Users get 10 minutes' free driving if the state of charge is below 20% when they return the EV to the charging station. We also use the prices of the control reserves from the transmission system operators of the corresponding regions (Transnet BW, Tennet, and California ISO).

We show that carsharing fleets can increase their gross profits without compromising the mobility of rental customers. This result is consistent across all three cities.

Profitability depends on whether there is an appropriate charging infrastructure and on the level of market demand for the control reserves that are available. We find that the largest profit increases for the VPP come from payments made for charging EVs when there is surplus energy that needs to be removed from the grid. Our data shows that the market rarely uses the EV batteries to cover electricity shortages, due to the high cost of batteries. However, both discharging and charging the EVs contributes to the bottom line of carsharing fleets. In the case of Stuttgart in particular, it is very profitable for the fleet; the profits from the VPP cover the cost of all the electricity used by their customers. Even with relatively low penetration of these vehicles at present, one of the benefits of the VPP that can immediately be realized is that carsharing fleets do not have to pay anything to fuel their EVs. In the future these cost reductions may be even more substantial as the demand for back-up power increases, due to the growing adoption of volatile, sustainable energy sources (Agricola, 2014).

Our research derives synergies by using EVs in the control reserve market instead of conventional, fossil-fuel-based means of peak-balancing, thereby creating greater efficiencies across the system. Renewable energy sources are weather-dependent and their production is difficult to forecast. Large-scale penetration of volatile energy sources poses a challenge to the stability of the grid (Kassakian and Schmalensee, 2011). The grid is the backbone of a highly perishable electricity supply chain, where supply and demand have to be in balance at all times. With the phasing out of power plants based on fossil fuels and a growing number of renewable energy sources, balancing the grid becomes increasingly difficult. In practice, this means that the chance of blackouts increases, with potentially disastrous consequences. Our research provides insights on how electric vehicles (among other energy sources) can be used to mitigate the instability of electric grids caused by the increasing amount of renewable resources. A key contribution of this research is that it increases the supply in the operating reserve markets, and therefore decreases the market price.

Mak et al. (2013) and Avci et al. (2015), who have been studying the location of battery-switching stations, consider future work on the area of intelligent charging of EV batteries as a crucial step in moving towards a sustainable economy. In the IS research community Watson et al. (2010) recognize the societal importance of the issue addressed in this paper by formulating an energy informatics framework with the aim of creating an ecologically sustainable society. Their framework formulates the need for IS research to take on the role of managing supply and demand in an energy-efficient way, and we show that we can do this by using EVs as mobile energy resources that are coupled by an information network that monitors the location, charging status, rental demand, and electricity supply and demand to create a real-time decision framework for optimizing resource utilization and profits. While forward looking, this topic is already receiving attention from the automotive industry. For example, Tesla Motors cars have a function to charge at cheaper night-time tariffs. With the Internet-of-Things the framework could be used not only in cars but also in individual appliances and devices (Porter and Heppelmann, 2014). An example is Google, which acquired Nest Labs with its programmable smart thermostats that produce energy savings of between 10% and 15% (Nest Labs, 2015). Creating appropriate IT infrastructure is central to the coordination mechanism for the current industrial examples as well as for the mechanism we propose.

To the best of our knowledge, this is the first study that uses real driving, charging, and locational data from more than 1000 EVs and makes an international comparison among three major cities in the USA, the Netherlands, and Germany with different energy mixes. Another key contribution of our research (from the perspective of EV balancing research) is that we assume that driving patterns are unknown a priori; this represents a key characteristic in EV balancing research, as previous work in this area by Vytelingum et al. (2011) and Tomic and Kempton (2007) was done using stationary batteries and EV fleets with known driving schedules respectively.

For the structure of our research and how it incorporates the design science aspects of relevance and rigor see Figure 1 (Hevner et al., 2004; Gregor and Hevner, 2013). In the Background and Related Literature Section we discuss the environment and knowledge base of our paper. Next, we give an overview of the data that is used to calibrate our model. This entails vehicle location, usage, and transaction data to calculate probabilities of rentals. Afterwards, we develop the bidding strategy, a design science artifact, in the Model Description Section. The challenge with this bidding strategy is that rentals are much more economically beneficial than using the battery for electricity storage, but it is wasteful if vehicles are idle. When building simulations to replicate real world phenomena one has to build a model that captures the essential characteristics of the environment without overfitting the data so that it is generalizable and applicable to other data. We describe the calibration of this simulation with bid and ask prices from real electricity markets in the U.S., Germany, and the Netherlands in the Evidence from a Real-world Setting Section. Consequently, we do a thorough evaluation and reflection of the bidding strategy artifact in the Analysis and Discussion Section. Finally, we conclude our paper in the Conclusions Section.

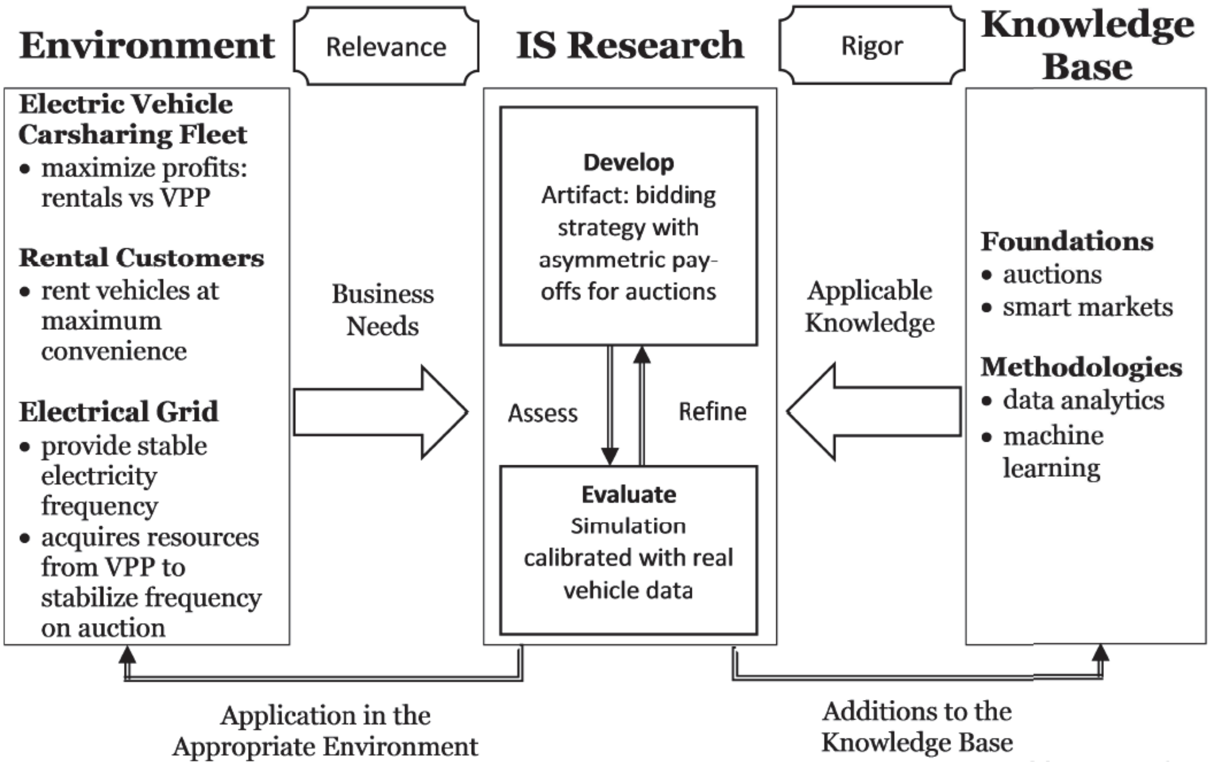

Figure 1: Information Systems Research Overview (Hevner et al., 2004) of Fleet-Power.

Background and Related Literature

This section summarizes relevant, previous research and outlines the general setting of balancing renewable energy sources. First, we describe the electricity market in detail explain how the trading prices are computed. Subsequently, we will position our work within the information systems literature on EVs, the carsharing context, and sustainability in general.

Balancing the Electrical Grid: Control Reserve Market

Electricity is sold on day-ahead and intraday markets as unit commitments hours before it is physically generated. However, when a source cannot meet its commitment (for example, due to technical problems or weather related issues) there are control reserve markets to guarantee immediate replacement (known in the US as the real-time market, and in the EU as the secondary control reserve market). These reserve markets require extremely fast reaction times called ramping rates from participating generators. For an overview of these markets and how they differ in their ramping rates, see Figure 2. EVs possess large electrical batteries whose energy is almost

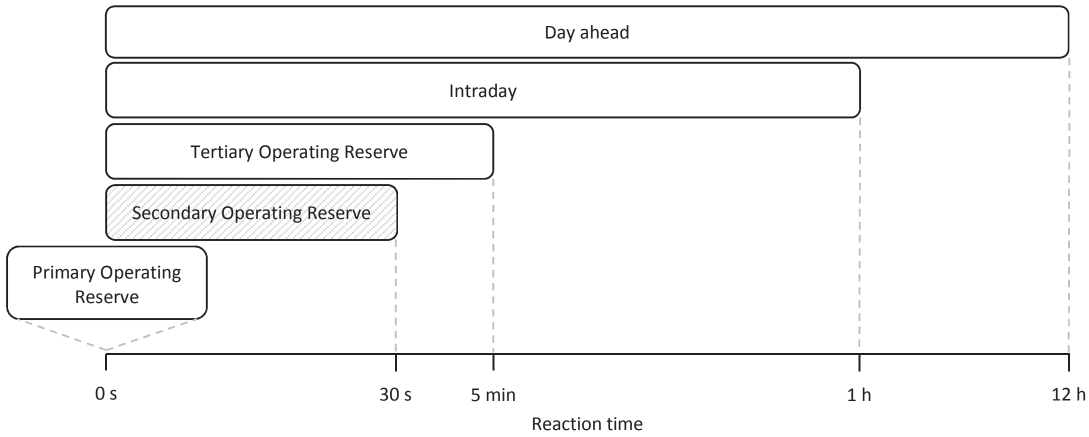

Figure 2: The reaction time, which is the time between notification and generation, differs according to the market. With no ramp-up time, EVs are able to comply with the short reaction times on the primary and secondary reserve markets.

instantly accessible without ramping cost, making them very suitable for reserve purposes. The present study focuses on the secondary control reserve market with a required ramp rate of 30 seconds (International Grid Control Cooperation, 2014). We focus on this market because the energy prices are higher than in markets which allow for a slightly longer ramp-up time. From this point on, when we are referring to energy markets we refer to the secondary control reserve market.

In the control reserve market, power plants are paid to be on standby so that they can produce (or consume) electricity when needed. The market is coordinated by electronic auctions, in which participants make *asks* or issue *bids*. The clearing mechanism is a multi-unit, first-price, sealed-bid auction, which is settled on a "pay-as-bid" basis (International Grid Control Cooperation, 2014). 'Asks' refer to the generation of electricity at short notice (*up regulation*), while 'bids' relate to the consumption of electricity, also at short notice (*down regulation*). Asks to generate electricity and bids to consume electricity state the price for which they would either generate or consume electricity and the maximum quantity they could generate or consume a week in advance. The transmission system operator settles these asks and bids as needed 30 seconds before delivery in merit order (the cheapest resources are used first).

We assume that in the future bids and asks can be placed separately for each 15-minute time interval. This future state is desired by TSOs to reduce entry barriers renewable energy sources to operating reserve markets (Agricola, 2014). This is a trend that is already see implemented partially today. Traditionally, the tender period

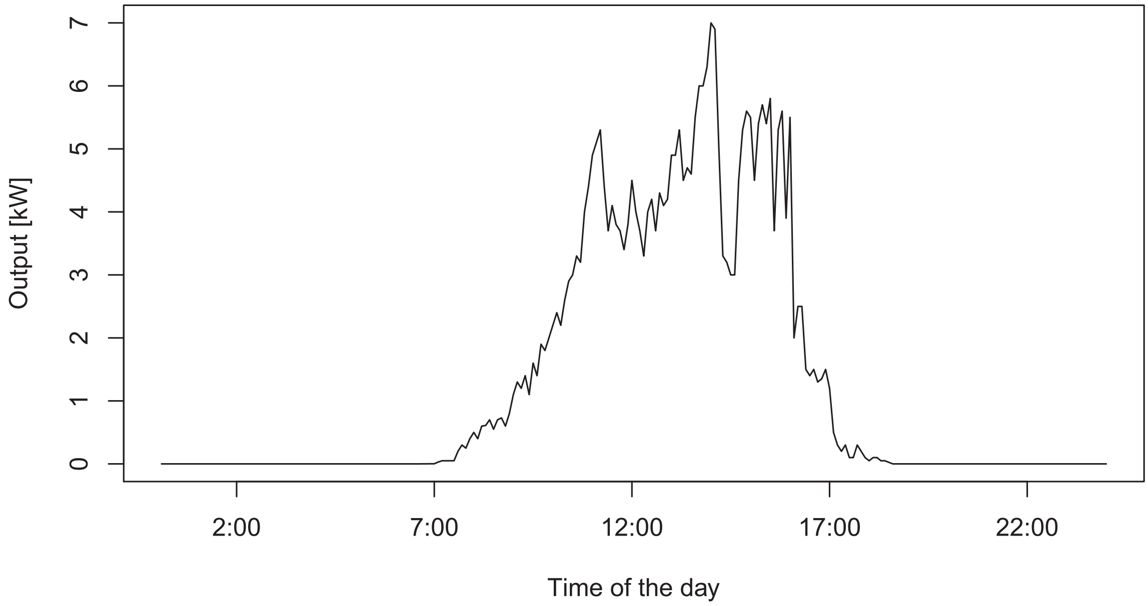

Figure 3: Solar panel electricity output illustrating the erratic behavior in photovoltaic electricity production with extreme variations in output. Data from Minneapolis, Minnesota in April 2015.

was monthly, up till now it was weekly, and from 2017 onwards it is going to be daily with 4 hour intervals. We are therefore confident that in the future 15 minute intervals are realistic given the trend to automated trading in the electricity sector. This market setup has implications for market power and competition. For a detailed discussion about this threat see Knaut et al. (2017).

Increasing levels of intermittent renewable energy and the decommissioning of conventional power plants exposes the control reserve market to the risk that at some point the demand may exceed the available supply (for instance, when the sun is suddenly covered by cloud and photovoltaic cells stop producing energy). Figure 3 provides an example of electricity output from solar panels, showing how production from solar panels is erratic, with extreme variations in output per minute. Note how the panel produces its maximum output at 1.30 pm, yet only minutes later production drops by more than 50% - in stark contrast to fossil fuel generators that produce electricity at a constant rate. These drops in energy output need to be offset within seconds to avoid blackouts.

Information-based Sustainable Society: Carsharing with Electric Vehicles

Information systems can be both a contributor to climate change and way of dealing with negative environmental impact. Similar to Loock et al. (2013)'s use of information to align individual interests with sustainability, we use information to align organizational goals with sustainability by means of a decision support system. As a result, financial and environmental goals are brought into harmony to foster carbon neutrality (Malhotra et al., 2013). Knowing when and where people rent EVs puts EV fleets in a position to make inferences about the rentals patterns of the population, and to make sociotechnical trade-offs (Geels, 2004) between their need for transportation and the need for storage on the energy market. We demonstrate this trade-off in a simulation platform similar to the one described by Ketter et al. (2016b,a) and calibrate it using real-world data.

Charging many EVs in the same neighborhood at the same time can quickly overload transformers and substations (Kim et al., 2012; Sioshansi, 2012). Previous research has addressed this issue by proposing smart charging, meaning that EVs are charged at times when the grid is less congested, helping to complement peaks in electricity consumption without creating new peaks. With smart charging EV fleets are given financial incentives to change their charging times, resulting in significant reductions in peaks (Valogianni et al., 2014). The departure times of EVs parked at public charging stations in California, Kara et al. (2015) show that an intelligent scheduling would result in a reduction of 24.8% in the monthly energy bill for users. An extension of smart charging is the V2G concept[1]. There are many successful models that use stationary storage to participate in electricity markets such as for example Mashhour and Moghaddas-Tafreshi (2011) or (He et al., 2016). Mashhour and Moghaddas-Tafreshi (2011) has built a model to bid in both energy and operating reserve markets. However, He et al. (2016) has done the same but included battery life cycles, which is an important consideration. We also consider battery life cycles, but additionally look at the unavailability of vehicles due to rentals, which adds an additional layer of complexity. A study by Vytelingum et al. (2011) considers the savings a household can make with a battery exposed to dynamic pricing on the energy wholesale market, and finds that efficient use of the battery would provide savings of 14% in utility costs and 7% in carbon emissions. Similar effects are found

[1]Vehicle-to-grid (V2G) discharging is technically possible. Even though not all charging stations support discharging yet, the standard of the International Electrotechnical Commission IEC 62196 supports V2G. For the purpose of this study, and with regard to future infrastructure, we assume that all charging stations have V2G capabilities.

by Zhou et al. (2015) in an industrial setting. Another study relating to EVs finds yearly benefits per EV of $176-203 (Schill, 2011). Tomic and Kempton (2007) show that the profitability depends on the target market: the larger the variations in the electricity price, the higher the profitability. Therefore, we focus on the control reserve market from Section .

Most studies make the assumption that households or car owners trade on the energy wholesale market. This assumption is not realistic, because they do not have a sufficient quantity of electricity to sell or buy to meet the minimum lot sizes required to participate in the market. To address this issue, Ketter et al. (2013) introduced the notion of electricity brokers (a.k.a. aggregators), which act on behalf of a group of households in order to reach the minimum lot size requirements. Simulations by Brandt et al. (2017) and Kahlen et al. (2014) show that this is possible to achieve with EVs. On top of the brokers we also apply the concept of a VPP. The asks and bids that are accepted constitute a promise to deliver electricity to the market, but which specific source will be used to fulfill that promise is not decided until actual delivery (i.e., whether a commitment will be delivered specifically from EV A or B (or a combination) is decided in real-time, based on the availability of those particular vehicles). We will show that this is a powerful tool which carsharing fleets can use to offer appropriate service levels for rental customers while making additional profits from balancing markets.

As the number of charging stations increases EVs are more likely to be connected to the grid and to be used as part of a VPP. As we want to make a statement about the profitability of VPPs of EV in the future, we also need to consider the possible density of charging infrastructures in the future. In earlier articles in this journal, Mak et al. (2013) and Avci et al. (2015) put forward an an optimal spatial infrastructure design for battery-swapping stations. This set-up has also been studied by Wolfson et al. (2011). However, we focus on conventional charging stations instead, because although there are 1,820 charging stations in Stuttgart, Amsterdam, and San Diego combined, there are no battery-swapping stations. We will therefore make recommendations on where additional charging stations should be placed.

A shortcoming of the existing studies is that - with the exception of Kara et al. (2015) - they all used either small fleets or data from combustion engine vehicles which have a longer range and are not subject to 'range anxiety', the fear becoming stranded with an empty battery. More importantly, in previous research, trips are assumed to be known in advance. In reality trips are more spontaneous (nondeterministic) and not always known in advance. This is problematic when an EV is committed to either charge or discharge at the same time as someone needs to drive it. Here a sociotechnical trade-off needs to be made between balancing the grid (technical) and providing mobility to customers (social). As it is impossible to determine precisely what value each individual places on mobility, we approximate its value with the profits from rental transactions. Free float carsharing, where users can pick up and drop off the vehicle anywhere allows us to specify the value for mobility for all pick up and drop off locations and times. Firnkorn and Müller (2011) show that free float carsharing has a significantly positive environmental effect, reducing carbon emissions by 6%. The optimal dimensions of the area where carsharing vehicles can be picked up and dropped off was studied by Wagner et al. (2016). For our study a free float carsharing business model fits very well as rentals are paid on a per-minute basis, and there still is uncertainty about where and for how long people will rent an EV (rentals are not booked in advance, i.e. they are nondeterministic), an issue which has not been covered in previous studies.

Previous research proposed the use of EV capacity as control reserves (Vytelingum et al., 2011; Schill, 2011), but it has not yet been shown whether this might be useful in a business setting. We show its usefulness not only in a realistic setting with 1,100 EVs, but also make an international comparison across the United States, the Netherlands, and Germany. Furthermore, our research can be generalized to all VPPs, and not just VPPs from EVs. Our model can be parametrized to represent different types of storage with different ramp rates. See Table 1 for an overview of the related literature.

Data

The fleet of Car2Go consists of 500 EVs in Stuttgart, as well as 300 EVs in both Amsterdam and San Diego. In addition to a sign-up fee, members pay for the carsharing service on a use basis only (per minute/hour/day, with an extra per km fee above a threshold of 50 km in Amsterdam and Stuttgart or 150 miles in San Diego). The prices across locations are similar but not identical (see Table 2), so we will use the arithmetic mean of the prices for ease of comparison between locations.

We chose Car2Go because it is a carsharing fleet with a global presence which uses the same cars (Smart ForTwo) across locations, thus allowing for a good comparison between countries. The sites are particularly suited for the purpose of this research because they are heterogeneous in terms of their energy mix. California and Germany

Table 1: Overview of related literature: Previous literature did not consider that driving patterns are not always known in advance and have not made an international comparison.

Author	**Real Data**	**Smart Charging**	**Vehicle-2-Grid**	**Virtual Power Plant**	**Charging Point Placement**	**Carsharing**	**Nondeterministic Driving Patterns**	**International Comparison**
Avci et al. (2015)					✓			
Brandt et al. (2017)	✓	✓	✓	✓				
Firnkorn and Müller (2011)	✓					✓		
Kahlen et al. (2014)		✓	✓	✓				
Kara et al. (2015)	✓	✓						
Ketter et al. (2013)				✓				
Kim et al. (2012)		✓	✓					
Mak et al. (2013)					✓			
Pudjianto et al. (2007)				✓				
Schill (2011)		✓	✓					
Sioshansi (2012)		✓						
Tomic and Kempton (2007)	✓	✓	✓					
Valogianni et al. (2014)		✓	✓					
Vytelingum et al. (2011)		✓	✓					
Wagner et al. (2016)	✓					✓		
Wolfson et al. (2011)					✓			
Zhou et al. (2015)		✓	✓					
FleetPower	✓	✓	✓	✓	✓	✓	✓	✓

Table 2: Environmental variables for Car2Go.

Description (variable, unit)[1]	**Stuttgart**			**Amsterdam**			**San Diego**		
EV type					Smart ForTwo Electric				
Max. battery capacity (Ω, kWh)					16.50				
Battery depreciation cost ($D, \frac{\$}{kWh}$)					0.13				
Charging speed per EV (γ,kW)			←		3.60 (linear)		→		
Discharging speed per EV (δ,kW)					3.60 (linear)				
Charging efficiency (ξ^{charge},%)					96.00 [2]				
Discharging efficiency ($\xi^{discharge}$,%)					97.40 (Reichert, 2010)				
EV fleet size (I)			491			343			367
Charging stations			1381			3561			541
Industrial electricity tariff ($ET, \frac{\$}{kWh}$)			0.12			0.10			0.10
Rental fee per minute ($)			0.37			0.39			0.41
Rental fee per hour ($)			18.92			18.92			14.99
Rental fee per day ($)			74.93			87.63			84.99
Extra fee after 50 km ($\frac{\$}{km}$)			0.37			0.39			0.45
	Min	**Avg**	**Max**	**Min**	**Avg**	**Max**	**Min**	**Avg**	**Max**
EVs available	148	395	491	64	253	343	106	281	367
EVs at charging station	32	104	232	21	70	122	2	34	64
Rental revenues (RB,$)	4	17	134	4	15	133	4	18	133

[1] Exchange rate is 1 Euro = 1.34 U.S. dollars.
[2] Value unknown, assumption based on our best estimate and talks with industry experts.

are both at the forefront of renewable energy penetration and standards, while the Netherlands rely on a more conservative electricity supply based on fossil fuels .

The rental and driving data was retrieved from a private application programming interface which we were given access to by Daimler. We retrieved a list of all EVs that were available for rental at the time of the query from the Car2Go website, `www.car2go.com`. We downloaded the data, added a time stamp, and stored it in a database every 15 minutes from May 1, 2014 to June 29, 2015; we were also continuing to collect the data for future research. This information contains the unique car name, the geographic coordinates of where the car is parked, the street name and zip code of that location (l), the state of charge of the battery (SoC), the state of the interior and exterior, and whether the EV is currently charging. We infer certain information about the transaction, such as how long the EV was rented, how many kilometers were driven, and how much revenue was earned as rental benefit ($\widehat{RB}$) by looking at the duration and timing of when the EV was unavailable for rent and the difference in the SoC level beforehand and afterwards. Even though the number of kilometers that can be covered using average fuel consumption will depend upon individual driving behavior, and this could therefore affect the accuracy of our estimates, we are confident that the differences will in fact be marginal, since all the journeys take place within the same urban environment. We assume that a fully charged EV will cover a distance of 66 miles (106 km). See Table 3 for an extract from the raw data and the information that we infer from it.

A drawback of the data set is that there is a chance that a car may be returned and rented again to another customer within the 15-minutes time interval. However, for the sake of our analysis, the EV remains unavailable, so this does not have a significant influence on the overall estimation and results. We also observe that several times particular EVs did not feature in the data for more than two days, even though the maximum rental duration is 2 days. We speculate that these cars were either in maintenance, repair, or not able to drive for some other reason, and were therefore not shown as available by Car2Go. We therefore removed from the dataset all rentals which we inferred from the data to have lasted more than two days. For a graphical illustration of how the rentals are distributed over a city see Figure 4, which shows the annual rental density in Amsterdam.

We infer the location of charging stations based on the GPS coordinates of where cars have been charged at least once in the dataset. We assume that if a car is parked at a charging station, it will be connected to the charging station. This is a sound

Table 3: This is an extract from the raw data. From this data we infer that a customer drove 7 miles (12km), rented the EV for 45 minutes, and paid \$14.99 (tariff for a full hour).

Car ID i	**Time** t	**State of charge** q	**Latitude**	**Longitude**	**Interior**	**Exterior**	**Street**	**Zip code** l	**City**	**Charging**	**Engine**
Car#1	12.05.2014 17:00	71%	32.76393	-117.122	good	good	Felton St 4728	92116	San Diego	false	electric
Car#1	12.05.2014 17:15	71%	32.76393	-117.122	good	good	Felton St 4728	92116	San Diego	false	electric
Car#1	12.05.2014 18:00	60%	32.76556	-117.168	good	unacceptable	Fashion Valley Rd 1261	92108	San Diego	false	electric
..						..					

assumption, because cars are only allowed to park at a charging station when they are plugged in, and any car that does not comply with this may be towed away.

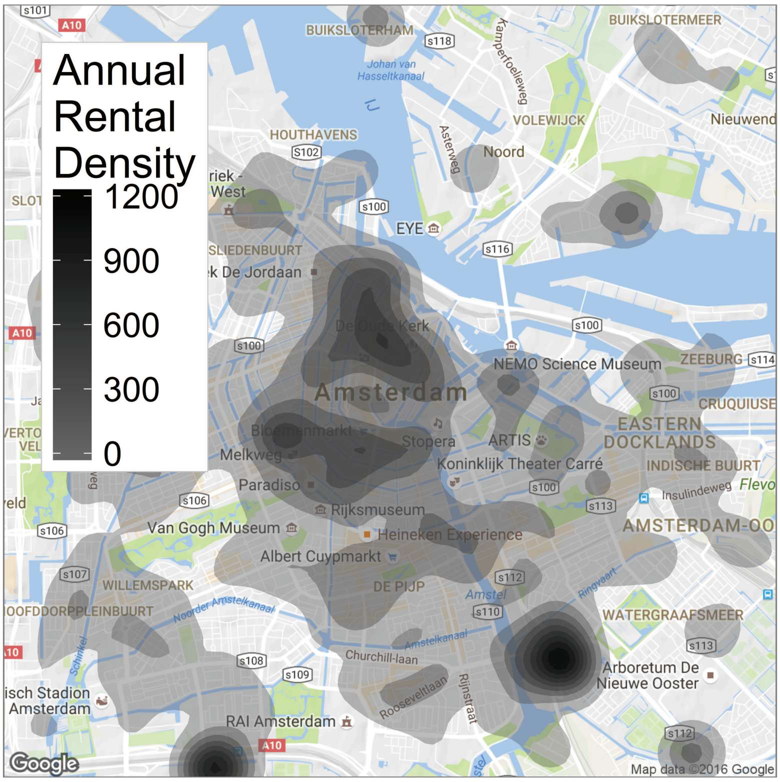

Figure 4: Shows where rentals occur most frequently in Amsterdam. The city center (middle), and business districts Zuidas (south west) and Amstel (south east) have particularly high densities of rental transactions.

Model Description

At the core of this research is the development of a decision support system that places bids and asks in a market. The market's clearing mechanism ultimately decides when to turn EVs into VPPs. The system is evaluated in a simulation environment bootstrapped with real-world carsharing driving data from Car2Go. A discrete event simulation is most suitable for this purpose, as we are dealing with a complex system that would be prohibitively expensive to build in the real world and where market parameters would be difficult to manipulate. In the following sections, we outline our selected approach and the reasoning behind it.

Virtual Power Plant Decision Support: FleetPower

Fleets need to decide how to deploy EVs by deciding which ones should be charged, which should provide V2G services, and which should be made available for rental, and they then bid accordingly. Figure 5 provides an example of how EVs are allocated in San Diego. The charging and discharging (V2G) is physically constrained to EVs that are connected to a charging station. Making real-time deployment decisions in this complex environment requires automated decision-making by an intelligent trading agent (Collins et al., 2010). We call this intelligent trading agent, which acts on behalf of the fleet, *FleetPower*. How FleetPower bids for the charging and discharging energy of the electric vehicles is described in the activity diagram in Figure 6. The agent needs to submit asks and bids for whatever price it is willing to charge or discharge and to decide how many EVs it wants to make available. The asks and bids need to be placed before the auction closes, and after the auction the agent has to provide or consume whatever quantity of energy has been agreed for any accepted bids or asks. The first step ❶ for the agent is to forecast the total amount of energy stored and how much can still be stored in the EVs that are available (parked at charging stations) for the timeslot under consideration. Next ❷, the agent has to determine a price at which it would be willing to sell or buy energy, to at least cover the opportunity cost. Afterwards ❸, this information about the price and the SoC of the EVs forms the basis of the asks and bids to be submitted to the auction. After the asks and bids have been submitted to the auction the market decides which asks and bids to accept and reject according to the pay-as-bid mechanism. The fleet needs to make sufficient EVs available to match the quantity of energy that has been agreed. ❹ These are dedicated EVs that deliver or consume energy according to the accepted asks and bids. ❺ If a customer asks to rent one of these particular EVs, either ❻ another car connected to a charging station then replaces that EV in the VPP in order to deliver the agreed amount to the market, or ❼ the customer is told that no car is available and ❽ the potential revenues are written off as opportunity cost. In practice customers will not be turned away as these cars will not show up on the list of available EVs, so customers would not notice any difference, especially since this is done already for cars that are charging. Note, that if another EV was free in the immediate vicinity, it was assumed that this car would then be rented out instead. Our interpretation of immediate vicinity is that customers are likely to be willing to walk to another car if it is approximately 250 meters away (drawn from a normal distribution with mean of 250 meters and standard deviation of 100 meters (d)). This value seems realistic to us, but we have also tested a mean of 100 and 500 meters with no significant difference in results. The great-circle distance between the coordinates is calculated using the haversine formula (Robusto, 1957). For a graphical illustration of the walking distance, see Figure 7. In the next section we will explain each step of the bidding procedure in more detail.

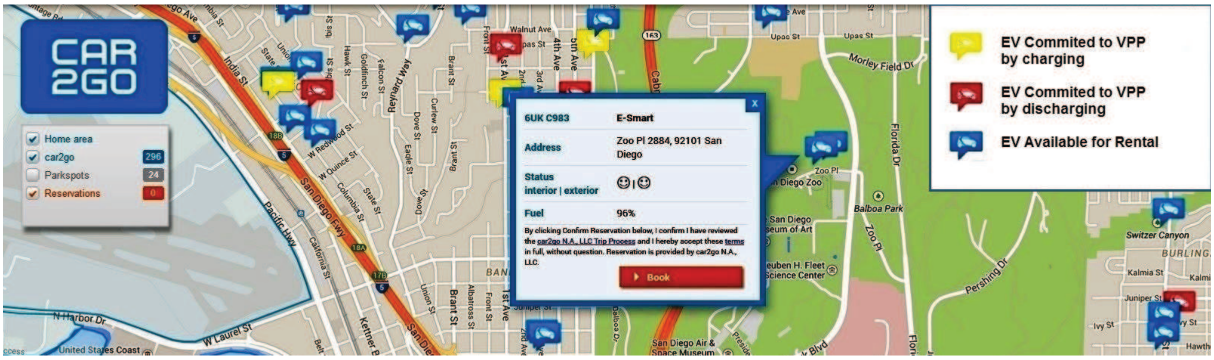

Figure 5: EV's from Car2Go in San Diego, USA. FleetPower committed strategically placed EV's as virtual power plants to charge or discharge (V2G).

Determine Ask and Bid Quantity ❶.

The first step in the ask and bid submission is to determine the quantity of electricity that should be offered in each 15-minute time interval. While it is important for customers to rent a car at a specific location, the precise location within a city is less relevant for energy markets as long as the car is parked at a charging station on the same distribution grid. Rather than making a decision on how each individual EV should be deployed, we can estimate an overall quantity of energy to charge and discharge, which allows us to harness the 'risk pooling effect'. This effect refers to the fact that EV storage potential and energy stored can be predicted more accurately for a whole fleet rather than for each individual EV. It is easier to make an accurate estimate of the number of cars that will be rented out in Amsterdam on a Sunday between 5.00 pm and 5.15 pm than to predict which specific EVs will be used.

For the purposes of this study, we have applied various machine learning algorithms, including neural network regression, support vector machine regression, and random forest regression, in order to forecast the energy storage available for charging (Q_t^{charge}) and discharging ($Q_t^{discharge}$) for the whole fleet at a specific time. We chose these regression algorithms because we use many attributes and there is a dependency in the data between the independent and dependent variables.

At the end of every week, the market closes for submissions for all the 15-minute intervals of the following week. We are therefore interested in predicting storage

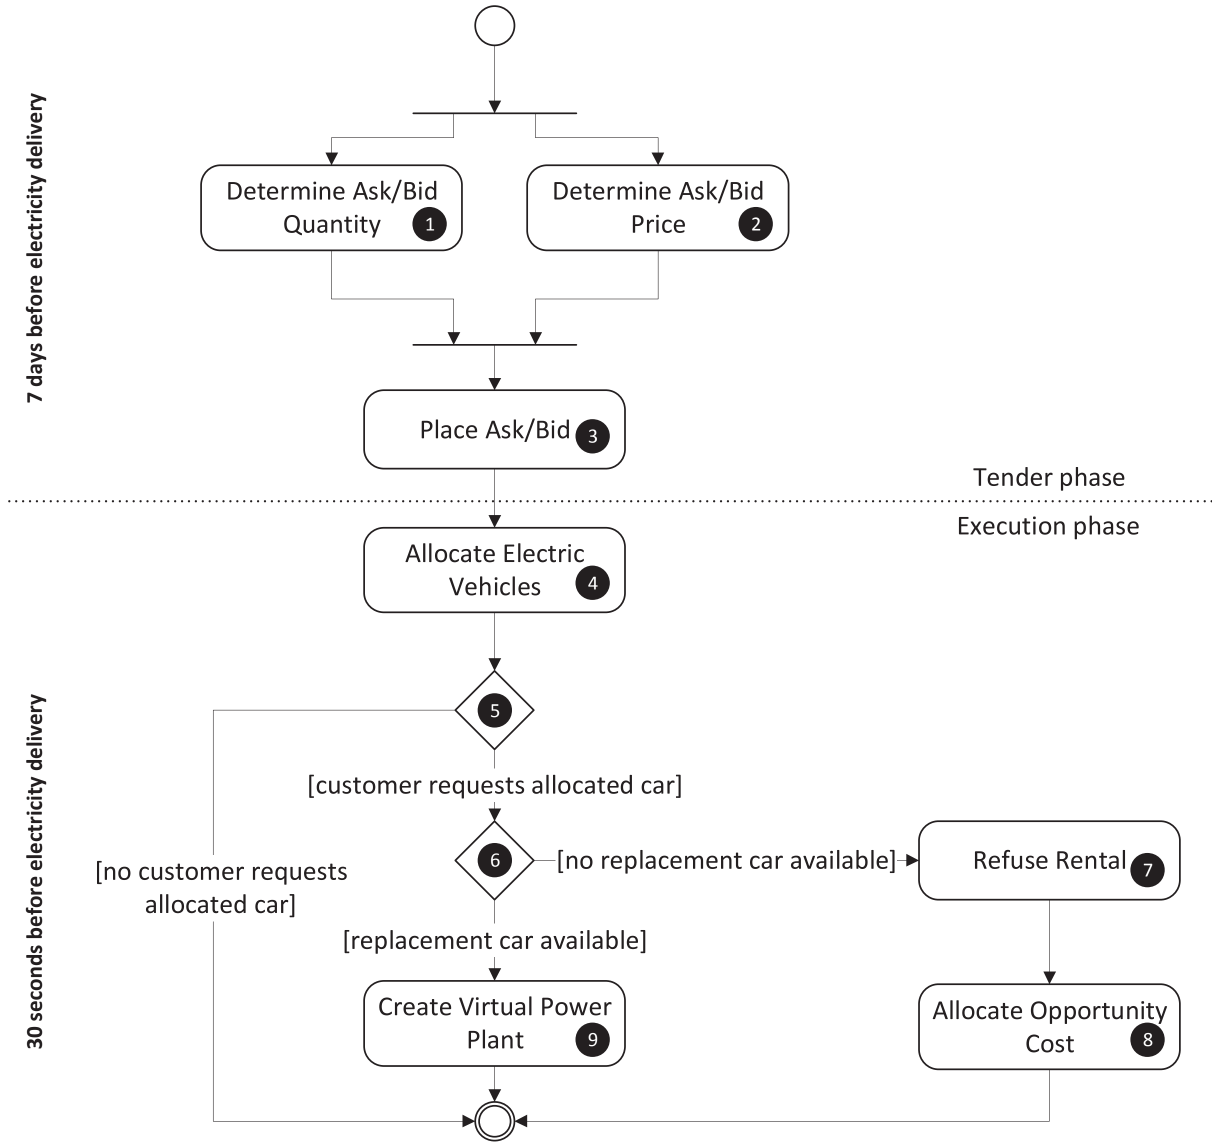

Figure 6: Activity diagram showing the decision-making steps involved in FleetPowers bidding on the secondary control reserve (real-time) market for energy.

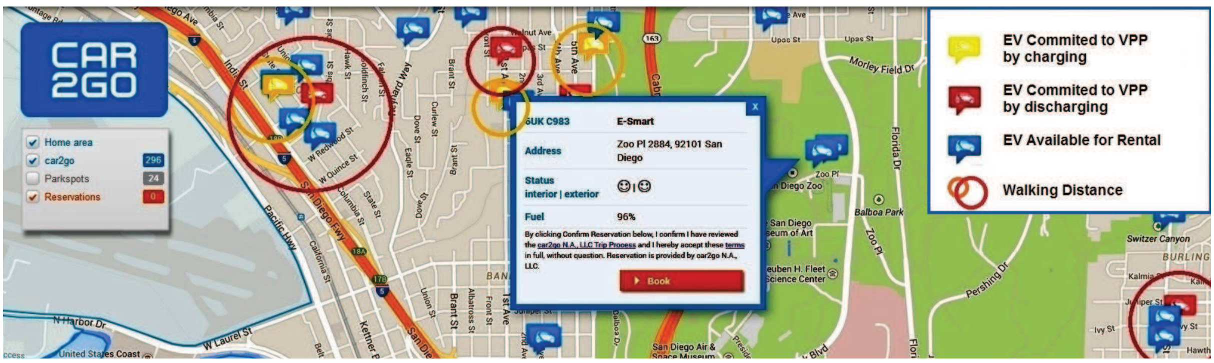

Figure 7: EVs from Car2Go in San Diego, USA. The circles indicate people's willingness to walk if the car they wanted in the first place is not available.

availability for up to one week in advance. The capacity to store or discharge for the fleet of EVs as a whole is predicted using the following equations:

$$Q_t^{charge} = \beta_{t,0} + \beta_{t,1} * day\_of\_week(t) + \beta_{t,2} * hour\_of\_day(t) \tag{1}$$

$$Q_t^{discharge} = \beta_{t,0} + \beta_{t,1} * day\_of\_week(t) + \beta_{t,2} * hour\_of\_day(t) \tag{2}$$

where β_0, β_1, β_2 are unknown parameters. The decisive factors determining the availability of storage are the day of the week (day_of_week), and the hour of the day in 15-minute intervals (hour_of_day). To predict the energy storage for each week we use a fixed two months time period (training period) to learn the daily and weekly patterns. The training period duration is fixed so that all predictions are comparable. The duration of the training period will be discussed in more detail in Section . A random forest regression model had the highest accuracy of prediction during the training period. This model was parametrized with two randomly preselected variables (mtry=2), 1000 randomized trees, and a minimum sum of weights for splitting of 5. We do not consider time series models, because commitments are due one week in advance. Special events such as soccer games, elections, or national holidays, and their influence on the model, will be discussed in Section .

In its essence the issue we are dealing with is a classification problem. We have to decide how many EVs we should assign to which class (rental or VPP). However, there is an asymmetric pay-off between assigning EVs to certain classes. Renting earns the carsharing fleet \$17.61 on average per transaction and the VPP earns the fleet \$0.09 on average per transaction. In addition, asks and bids on the energy market are binding and non-delivery will result in very high penalties. We therefore give misclassifications for the rental class proportionally more weight than the VPP class. This weight is assigned with the stratified sampling method, where we sample disproportionately to reflect the asymmetric pay-off (Berk et al., 2009). This method decreases the likelihood that our model adds a car to a VPP.

Example: Assume, for instance, that we are interested in submitting a bid in Amsterdam for Sunday, July 6, 2014 for the time interval t 5.00 pm to 5.15 pm (t=60 as it is the 60^{th} 15-minute interval). To predict the available storage we look at the training period from May 1, 2014 to the day on which asks and bids can be submitted for auction on June 30, 2014. Based on the number of EVs and their state of charge for each Sunday in that time period, as well as each t=60 time period, we predict the availability for the test period July 6, 2014 at t=60. To account for changes in usage patterns over time we explicitly include the availability between t=60 of the last Sunday as a lagged dependent variable (in this case: June 29, 2014). With more historical data one could also include the same day from previous years to improve the accuracy of the model. If there were on average 10 EVs connected to charging stations, each with a state of charge of 70% (SoC) and a 16.5kWh battery (Ω), the storage available for charging would be $Q_t^{charge} = 10 * 0.3 * 16.5kWh = 99kWh$, and the storage available for discharging would be $Q_t^{discharge} = 10 * 0.7 * 16.5kWh = 231kWh$. Due to physical constraints of the available infrastructure in Stuttgart, Amsterdam, and San Diego, the charging (γ) and discharging speed (δ) of 3.6kWh per hour (or 0.9kWh in 15 minutes) and charging (ξ^{charge}) and discharging efficiencies ($\xi^{discharge}$) of 96% and 97.4% limits the actual values to a maximum of $10 * 0.9 * 0.96 = 8.6kWh$ and $10 * 0.9 * 0.974 = 8.8kWh$, rather than 99kWh and 231kWh for the 15-minute interval.

Determine Ask and Bid Price ❷.

The second step in the ask and bid submission process is to determine the price at which the asks and bids should be offered so as to balance out potential gains to be made from the auction versus the likelihood of the offer being accepted. There is a price for capacity (standby fee) and for electricity (per unit of energy). We bid at a capacity price of $0\frac{\$}{MW}$ to ensure that our bids will always be considered by the System Operator.This also makes the strategy independent of whether or not a capacity price will be applied in the future with shorter time intervals. We show that the bidding strategy at $0\frac{\$}{MW}$ is economical even if there is a capacity price component, but the claim of optimality does not apply if there is a capacity fee. To determine the electricity price we apply a bottom-up model which estimates the optimal price per EV. For each EV, the fleet has a number of costs that need to be covered. For example, when charging an EV, the agent needs to ask a price P^{charge} that takes into account the alternative, which would be simply paying the industrial electricity tariff, the opportunity cost of being unable to serve a customer while charging, plus a margin. To discharge the EV (V2G), an agent should ask a price $P^{discharge}$ that is based on the energy cost of charging in the first place, the cost of battery depreciation, the opportunity cost of not being able to serve a customer while discharging, and of not being able to serve customers in the future due to a lower battery SoC, plus a margin.

For a table of notation, including measurement units, see Table 4.

Determine the price for charging (P^{charge}): EVs can be parked anywhere in the city, but only if an EV is parked at a charging station does FleetPower have

Table 4: Table of notation.

Variable	Description	Unit
$\widehat{RB}$	Expected rental benefit per unit of energy stored	$\frac{\$}{kWh}$
RB	Observed rental benefits	$
C^{charge}	Charging cost, see Equation 3	$
$C^{discharge}$	Discharging cost, see Equation 6	$
c	Rental probability, see Figure 9	$\frac{\%}{100}$
D	Battery depreciation cost	$\frac{\$}{kWh}$
d	Distance EV_i to closest EV available for rent	km
EC	Energy cost, based on P_t^{charge}	$\frac{\$}{kWh}$
ET	Industrial electricity tariff (flat price)	$\frac{\$}{kWh}$
I	Total number of EVs	
i	Specific EV	ID
l	Location	zip code
P	Bid/ask price for buying or selling electricity from reserve market	$\frac{\$}{kWh}$
Q	Bid/ask quantity for buying or selling electricity from reserve market	kWh
Q^*	Equilibrium quantity (sign indicates shortage or surplus electricity)	MWh
q	State of charge (SoC) (Ψ/Ω)	$\frac{\%}{100}$
t	Time interval	index
β	Unknown regression parameter	-
Δt	Duration of a time interval	0.25 hours
γ	Charging speed	kW
δ	Discharging speed	kW
λ	Dummy to account for opportunity costs from recharging	boolean vector
μ	Margin on the bid/ask price, to optimize the bidding price	$\frac{\$}{kWh}$
ξ^{charge}	Charging efficiency	$\frac{\%}{100}$
$\xi^{discharge}$	Discharging efficiency	$\frac{\%}{100}$
Ψ	Amount of electricity stored in an EV	kWh
Ω	Maximum battery capacity	kWh

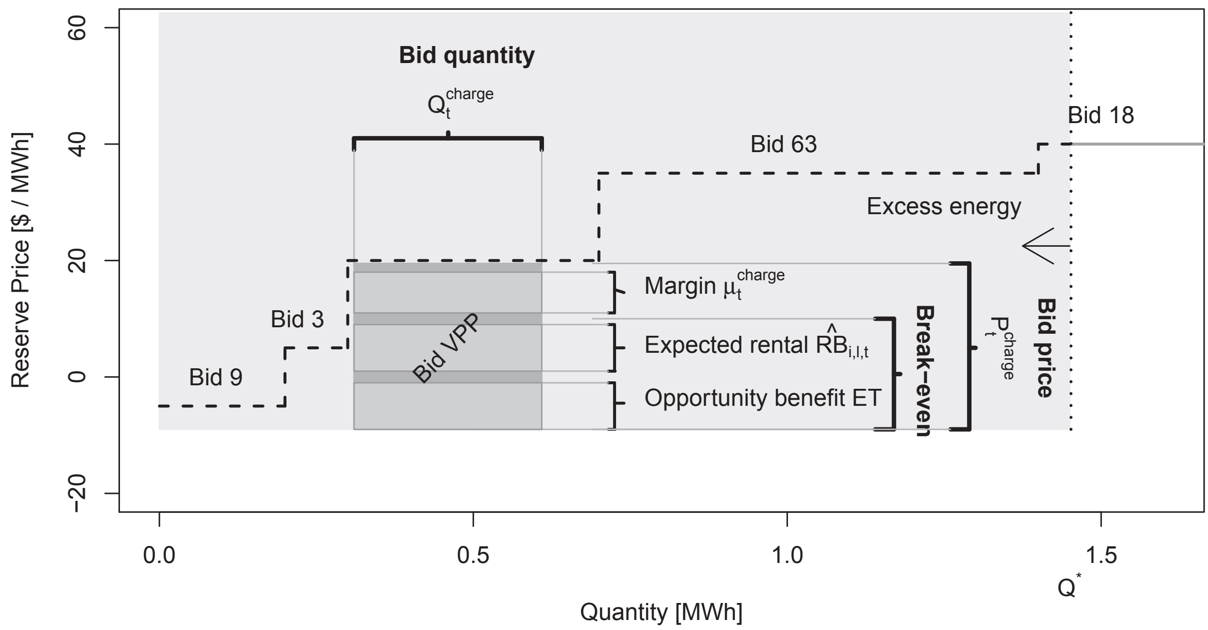

Figure 8: The graph shows the demand curve for getting rid of surplus energy; in this example the grid needs to get rid of 1.45 MWh. As our VPP bid offers to consume this electricity more cheaply than bid 18, which is the last (partially) cleared bid, the market agrees to sell 0.4 MWh to the fleet.

the option to turn it into a VPP. Where this is the case, FleetPower can bid in the energy market for a cheap electricity rate. Bids are submitted to the auction, and Figure 6 shows the components of the bid, together with the bid quantity and price. The first component of the bid price, opportunity benefit, serves as a reference point; FleetPower will not purchase electricity on the energy market if the industrial electricity tariff were cheaper. The second component, the expected gross profits from rental, ensures that EVs are less likely to charge when it is probable that they will be rented out. The industrial electricity tariff, and the expected rental profit determine the break-even point at which renting out or turning an EV into a VPP are equally matched financially. The final consideration, the margin, allows the fleet to make a profit, and here a trade-off needs to be made between the pay-off and the likelihood of the bid being accepted. This margin parameter optimizes the profits for the fleet owner based on the probability of being called upon in the training period to maximize profits. We now describe the bid in more detail.

The financial cost of charging (C^{charge}) a specific EV i at 15-minute time interval t is determined by the following equation:

$$C_{i,l,t}^{charge} = \frac{\min(SoC_{i,l,t} * \Omega, \delta * \Delta t)}{\xi^{charge}} \left(-P_{i,l,t}^{charge}\right) \tag{3}$$

where $min(SoC_{i,l,t} * \Omega, \delta * \Delta t)$ is the amount of electricity that could still be charged to the battery of car i at interval t and (P^{charge}) is the bid price to charge, which differs per EV i and time interval t. The variable ξ^{charge} accounts for the charging inefficiency.

The bidding price P^{charge} for charging is determined as follows:

$$P_{i,l,t}^{charge} = ET - \widehat{RB}_{i,l,t} - \mu_t^{charge} \tag{4}$$

where ET is the opportunity benefit of not having to pay the industrial electricity tariff, $\widehat{RB}$ is the expected rental benefit which we will describe next, and μ^{charge} is the profit margin which is parametrized to maximize the overall profits for all previous time intervals t (of the training data set). In other words, we take the electricity tariff, we deduct what we could have earned with the EV if it had been available during that period, and we add a margin to arrive at the lower electricity price the carsharing fleet would be willing to accept in return.

We use the same machine learning algorithms to predict the rental profits per unit of energy and thus to decide how much energy to offer (see ❶ in Figure 6). In contrast to the quantity prediction, support vector machine regression had the best predictive accuracy for rental profits (see ❷ in Figure 6). The rental profits are predicted with the following equation, similar to Equations 1 and 2:

$$\widehat{RB}_t = \beta_{t,0} + \beta_{t,1} * day\_of\_week(t) + \beta_{t,2} * hour\_of\_day(t) \tag{5}$$

The support vector machine regression was parametrized using a radial basis kernel function with the parameters gamma = 2 and cost = 1. The expected profit for renting EV i per unit of energy ($\widehat{RB}$) during interval t parked at location l is determined by a support vector machine regression with four independent variables: rental probability c, state of charge of the battery SoC, interior status, and exterior status.

The rental probability c captures the preferences and behaviors of those who rent EVs. The probability has three dimensions: location, the hour of the day, and the day of the week. The locational dimension gives insights into the likelihood that EVs are rented out given that they are parked in a certain district, which is represented by a zip code. The temporal dimension gives insights into the likelihood that EVs will be rented out given a certain hour of the day and a day of the week. We break time down into discrete 15-minute intervals for each day of the week. Based on this information, we create a four-dimensional model that enables us to predict whether a car is likely to be rented out within the next 15-minute interval, based on the day of

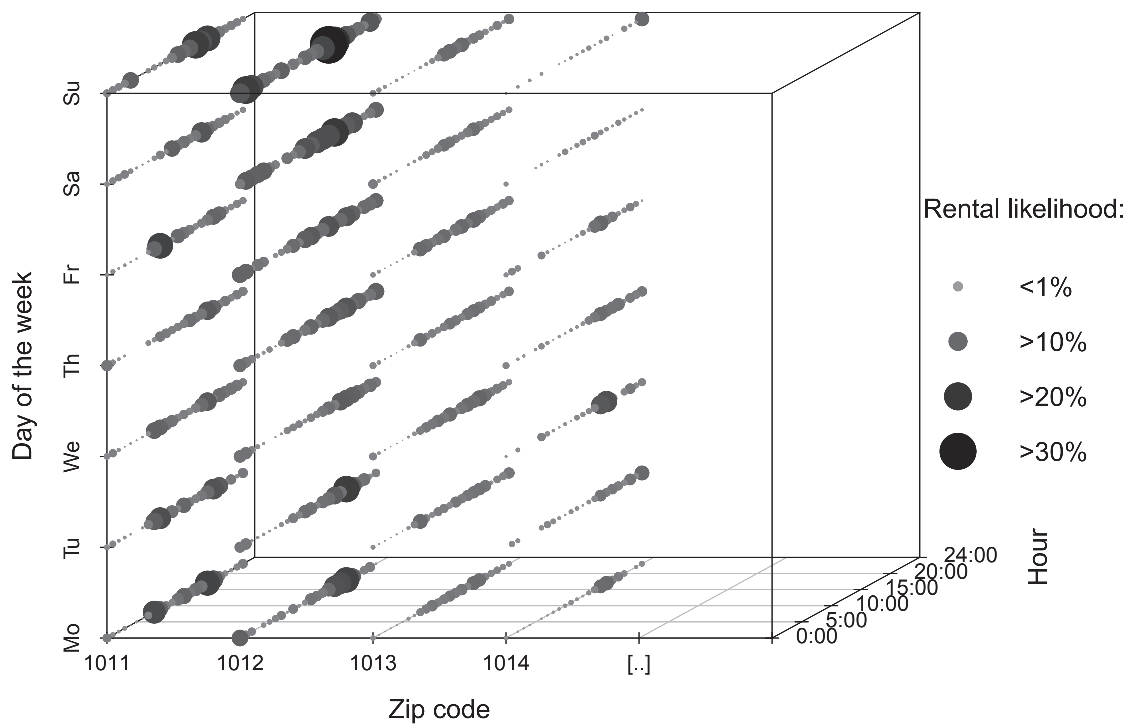

Figure 9: Four-dimensional rental likelihood model that shows when EVs are rented out in Amsterdam, based on the day of the week, hour of the day (15-minute interval), and zip code of where they are parked. There is a distinct difference in rental likelihoods - for instance, between the central station (1012) and the periphery (1014).

the week, the specific 15-minute interval, and the zip code of the place where the car is parked. Unlike day and time, zip code is a categorical variable. Figure 9 provides an example from the city of Amsterdam, which serves as an input to the support vector machine regression model used to predict the rental transactions. There are significant differences between locations: for example, on a Sunday morning in Amsterdam in zip code 1012 (central train station) the rental likelihood of 5 to 10% is much higher than in zip code 1014 (periphery), where the likelihood is below 1% .

Example: How we determine the bid price for charging is illustrated by what happens with car $ID\#5$ on the morning of Monday May 11, 2015 between 5.15 am and 5.30 am. The battery is charged (SoC) to 90%, the interior and exterior status are "good", and the car is parked in zip code (l) 1012 (Amsterdam central station). At this time and location the likelihood of the car being rented out is below 0.3%, and the trained support vector machine regression over the training period (60 days: April 11, 2015 to May 11, 2015) predicts an expected rental benefit of $0.042\frac{\$}{kWh}$ (over the kWh left in the battery). Given an industrial electricity tariff of $0.08\frac{\$}{kWh}$ and an optimal margin of $0.02\frac{\$}{kWh}$, the bidding price is $P^{charge}_{5,1012,21} =$

$0.08\frac{\$}{kWh} - 0.042\frac{\$}{kWh} - 0.02\frac{\$}{kWh} = 0.018\frac{\$}{kWh}$. If the quantity $Q^{charge}_{5,1012,60} = 0.9kWh$ is bought from the market, and an adjustment is made for efficiency losses ξ^{charge}=0.02, the total opportunity cost of charging the EV during that period would be $C^{charge}_{5,1012,60} = (0.9kWh/(0.98)) * (0.08\frac{\$}{kWh} - 0.018\frac{\$}{kWh}) = \$0.057$. In order for the charging of this EV to be economical during this 15-minute interval, the carsharing fleet would need to bid at a price not exceeding \$0.057. If it would pay more for electricity it would be better off to use the flat electricity tariff. Any figure above this would mean that the fleet would be better off charging its vehicles using electricity supplied at the standard flat tariff.

Determine the price for discharging/V2G ($P^{discharge}$): An EV can also contribute to a VPP by discharging if it is parked at a charging station with V2G. FleetPower then has the option to sell electricity through V2G by submitting an ask to the energymarket. Figure 10 shows an ask quantity and ask price, broken down into its various components, and compares these to other asks in the auction. The first component of the ask price (energy cost) is the sum needed to reimburse the fleet for the cost of charging the EV in the first place. The second component, battery depreciation, compensates the fleet for wear and tear on the battery. The third component, the expected rental gross profit, ensures that the maximization takes into account that an EV is less likely to discharge using V2G when it is probable that it will be rented out, and this calculation includes an allowance for the time needed to recharge the EV to its previous charge state. Even though rental gross profits also include an element to cover the costs of battery depreciation, we explicitly include this as a separate part of the ask price, because there are substantial differences between discharging and driving in terms of the battery depreciation depending on the volume of activity. For example, at night the expected battery depreciation costs from rentals are close to zero, because it is unlikely that someone would rent an EV, whereas if the ask is accepted, the battery depreciation costs associated with discharging and subsequent recharging will be incurred in full. Also, in this case the electricity cost, the battery depreciation cost, and the expected rental profits are combined to determine the break-even price at which the rental and VPP are of equal value to the carsharing fleet. The last consideration, the margin, allows the fleet to make a profit in the "pay-as-bid" market, though a trade-off needs to be made between the potential gains and the likelihood of the ask being accepted. We describe the ask in more detail below.

The financial cost of discharging ($C^{discharge}$) a specific EV i at 15-minute time interval t is determined by the following equation:

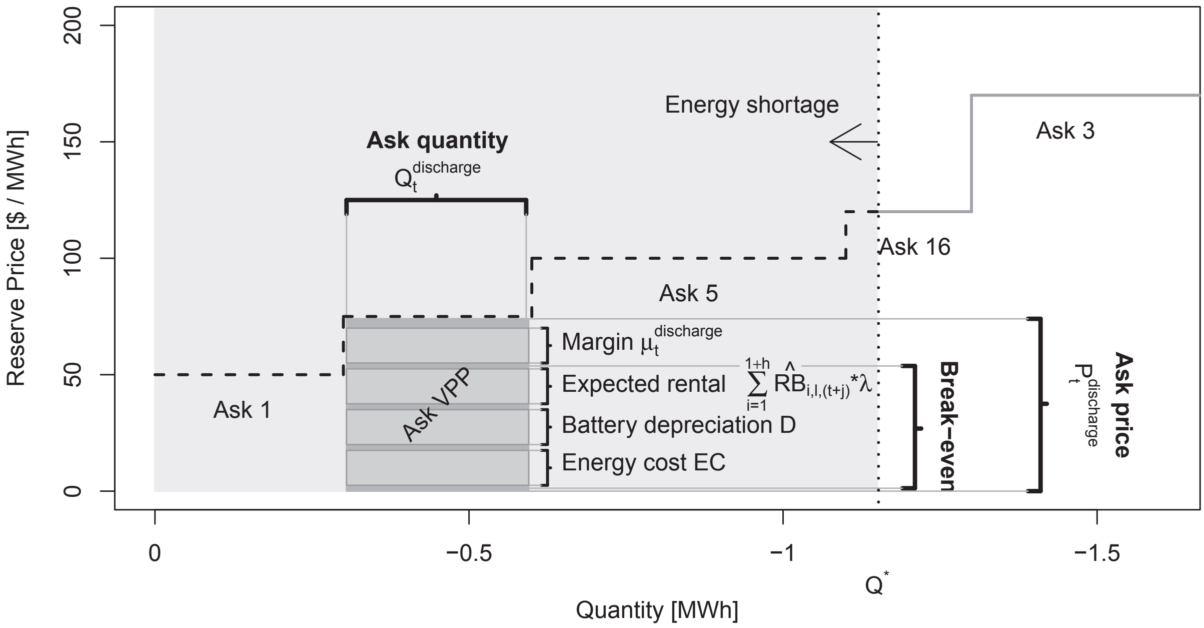

Figure 10: The graph shows the supply curve for the purchase of energy to bridge a deficit in the grid; in this example, an additional 1.2 MWh is required. As our VPP ask offers to provide this electricity more cheaply than ask 16, which is the last (partially) cleared ask, the market agrees to buy 0.3 MWh from the fleet.

$$C_{i,l,t}^{discharge} = \frac{Q_{i,l,t}^{discharge}}{1 - \xi^{discharge}} \left[-P_{i,l,t}^{discharge} \right] \tag{6}$$

where ($Q^{discharge}$) is the electricity stored in EV i that can be accessed within time interval t. $P^{discharge}$ is the price at which the electricity is being offered for sale, as defined in Equation 7. $\xi^{discharge}$ is the discharging inefficiency that accounts for energy conversion losses.

The asking price for discharge ($P^{discharge}$) is determined as follows:

$$P_{i,l,t}^{discharge} = -D - EC - \sum_{j=1}^{1+h} ((\widehat{RB}_{i,l,(t+j)} * \lambda_{i,l,(t+j)}) - \mu_t^{discharge} \tag{7}$$

where the cost for wear on the battery is depreciated (D) for each kWh of energy used. Also, the energy costs for charging EC, based on the asks accepted during the training period, are taken into account. The summation term refers to recharging the EV after V2G. $h = \text{round}\left(\frac{Q_{i,l,t}^{discharge}}{\delta * \Delta t}\right)$ is the time it takes to recharge the EV, rounded to the nearest time interval (15 minutes). $\widehat{RB}$ are the opportunity costs of not being able to rent out the EV due to it being committed to a VPP during the current interval t and costs of recharging it subsequently. The dummy variable λ states that opportunity costs only apply if the next person to rent the vehicle cannot complete the expected trip $\widehat{RB}_{i,l,t+j}$ with the remaining capacity from V2G. $\mu^{discharge}$ is the margin which maximizes the overall profits for the time intervals t in the training data set in a similar way to the margin in Equation 4.

Example: To see how we determine the bid price for discharging, take the example of car $ID\#5$ on the morning of Monday May 11, 2015 from 5.15 am to 5.30 am at zip code 1012. The same conditions apply as in the example provided above for charging. Assume that the battery depreciation are $D = 0.1\frac{\$}{kWh}$, the quantity in question is $Q_{5,1012,60}^{discharge} = 0.9kWh$, the discharging speed $\delta = 3.6kW$ per EV, the rental benefit for the next time interval (5.30-5.45 am) is $RB_{5,l,61} = 0.055\frac{\$}{kWh}$, and $\lambda_{5,l,61} = 1$ as there are rental costs in t=60, but as the battery SoC is completely full it is unlikely that in period t=61 the EV will have too little battery power left to be used for another rental. Under these circumstances the price for discharging can be expressed as follows: $P_{5,1012,60}^{discharge} = -0.1\frac{\$}{kWh} - 0.08\frac{\$}{kWh} - \sum_{j=1}^{2}((\widehat{RB}_{5,l,(60+j)} * \lambda) - 0.02\frac{\$}{kWh} = 0.018\frac{\$}{kWh}$. If the quantity $Q_{5,1012,60}^{discharge} = 0.9kWh$ is bought from the market, and an adjustment is made for efficiency losses $\eta^{discharge}$=0.98, the total opportunity cost of discharging the EV during that time period would be $P_{5,1012,60}^{discharge} = (0.9kWh/(0.98)) * ((0.1\frac{\$}{kWh} +$

$0.08\frac{\$}{kWh}) - 0.018\frac{\$}{kWh}) = \$0.149$. In order for the discharging of this EV to be economical, the fleet would need to ask a price of at least \$0.149.

Place Ask and Bid ❸.

The third step is to combine the quantities and prices as asks and bids respectively and submit them to the market. To do this the agent chooses the EVs i with the lowest cost for charging $C_{i,l,t}^{charge}$ and discharging $C_{i,l,t}^{discharge}$ until the respective overall quantities Q_t^{charge} and $Q_t^{discharge}$ are reached. Each quantity is submitted to the energy market at the average price from Equations 4 and 7, weighted by the amounts bought or sold. We only submit one ask and one bid for each time interval due to the minimum lot size of 1MW. We do not consider submitting multiple asks and bids for the same auction, even though this would increase the profits, because substantially larger fleets would be required to meet the minimum lot size. To reach the 1MW threshold one would need to collaborate with an aggregator.

Example: Take, for example, the following situation where the cost of charging EVs ID#1 and ID#2 are $C_{1,1012,59}^{charge} = 0.036\$$ and $C_{2,1012,59}^{charge} = 0.09\$$, and the corresponding bidding prices are $P_{1,1012,59} = 0.04\$$ and $P_{2,1012,59} = -0.02\$$ (t=59 means 4.45-5.00 pm). The negative price for EV ID#2 means in the time interval t=59 the market needs to pay Car2Go for the charging to be economically worthwhile. The state of charge of the batteries of the EVs are $SoC_1 = 0.3$ and $SoC_2 = 0.4$ respectively. FleetPower has determined the optimal quantity that should be offered to the market to be $Q_t^{charge} = 1.5kWh$. We also assume a battery capacity of $\Omega = 16.5kWh$, and a charging speed (γ) of 3.6kWh per hour (or 0.9kWh in 15 minutes) per EV. In this case FleetPower offers to provide $1.5kWh$ at a price of $0.016\frac{\$}{kWh}$, as $0.9kWh$ (depending on the amount that can be discharged with the infrastructure in the time constraint $0.9kWh$ and what SoC the battery is in, $0.6 * 16.5kWh$) can be provided from EV with ID#1, which has the lowest cost, and the remaining 0.6kWh will be provided from EV ID#2.

Endogeneity from Market Participation.

By participating in the market we may have an influence on market equilibrium, and this might in turn lead other market participants to behave differently. However, we argue that there is no endogeneity problem from reactions to our market participation as the asks and bids of other participants are aligned with their preferences. Discriminatory-price multi-unit auctions are not incentive-compatible but our approach will work with any mechanism. For example, the uniform-price multi-unit auction can be designed to be posterior regret-free (i.e., even though the mechanism is not incentive-compatible a priori, no one could benefit from not bidding their true valuation when evaluating allocation ex-post) (Bapna et al., 2005). Under these mechanisms other market participants have no incentive to alter their behavior in response to new market entrants. Our methodology will also work well with this kind of mechanism. While the revenues may be different, the structural results will not change.

Evidence from a Real-world Setting

For the evaluation we consider the 14 month period from May 1, 2014 to June 29, 2015. We train our model from the first two months from May 1, 2014 till June 29, 2014 and test it on the first week of auctions (the bids and asks for all the 15-minute intervals in a week are always submitted for the full week in advance, Monday to Sunday). Consequently we use a rolling time window for the training period of two months for each week of bidding. We test the algorithm for each week in the period from June 30, 2014 to June 29, 2015. From these training sets values for the rental likelihood model, expected driven kilometers, rental time, and rental profits are used to train the model. Based on this training period we evaluate the trained model over all one-week bidding blocks in that time period. There is no need to simulate the distribution of trips as we have an immense number of real driving transactions with which we can test our results, and a calibration of the driving data is therefore not necessary (as it is real data). The test period is given externally by the market while the training period is a constraint from the data collection perspective (we are limited to 14 month of collected data).

Energy Market Data: California ISO, Transnet BW, and Tennet

As illustrated in the Section , EV storage is particularly suited to real-time market operation due to the fast response times required (dispatch occurs within seconds of order acceptance). We therefore use auction data from these markets to determine the prices for balancing (charging as well as discharging) at each point in time. We use the data from the energy market operators in Stuttgart, Amsterdam, and San Diego. In Stuttgart, we use data from regelleistung.net, the German energy market operator, in Amsterdam we use data from Tennet, and in San Diego the data comes from California ISO. The data for Stuttgart contains the individual bids and asks with the respective quantities and prices for each 15-minute time interval. From these bids and asks we form the demand and supply curves. The clearing point Q^* sets the equilibrium, which determines whether the energy market operator settles the asks and bids placed by FleetPower (if the price P from the model is below the market price). For San Diego and Amsterdam we only know the clearing prices, though this still allows us to infer which bids and asks are accepted (the ones below the market price). The violin plot in Figure 11 shows the average regulation prices and their standard deviations for all three cities. As can be seen in all three, the prices for regulation reserves are quite variable (standard deviation ranging from 18.2 to 94.3). The low renewable energy content in the Dutch energy mix is responsible for significantly lower discharging prices (41.6 compared to 68.2) and standard deviations than in Stuttgart. This reduces the revenues for VPPs in the Netherlands. In contrast, the high renewable energy content in the energy mix in Germany and California leads to generally higher prices for discharging in Stuttgart and to large fluctuations in price in San Diego. This increases the revenues for VPPs in those countries, because those in Stuttgart can charge cheaply and sell large volumes of energy at higher prices, or in in San Diego they can sell smaller amounts of energy at extremely high prices at the high variation in the evening hours.

Analysis and Discussion

In this section we will discuss the evaluation of the business model in terms of the profits to the carsharing fleets and the implications for the grid. We will also describe the sensitivity analysis that we have conducted to show the robustness of our model.

Accuracy of the Model and Economic Sustainability

The proposed model creates a new business model for carsharing fleets, which is a natural extension to the traditional carsharing business. We are interested in whether this VPP business can increase the profits of carsharing fleets. Specifically we are interested how VPPs influence the gross (variable) profits, excluding overhead cost, because they are incurred regardless of the VPP model. Figure 12 shows the gross profits of Car2Go in Stuttgart, Amsterdam, and San Diego, both for carsharing only and for carsharing combined with a VPP. While the gross profits (black line) are volatile throughout the year, there is a clear seasonal pattern in the winter period

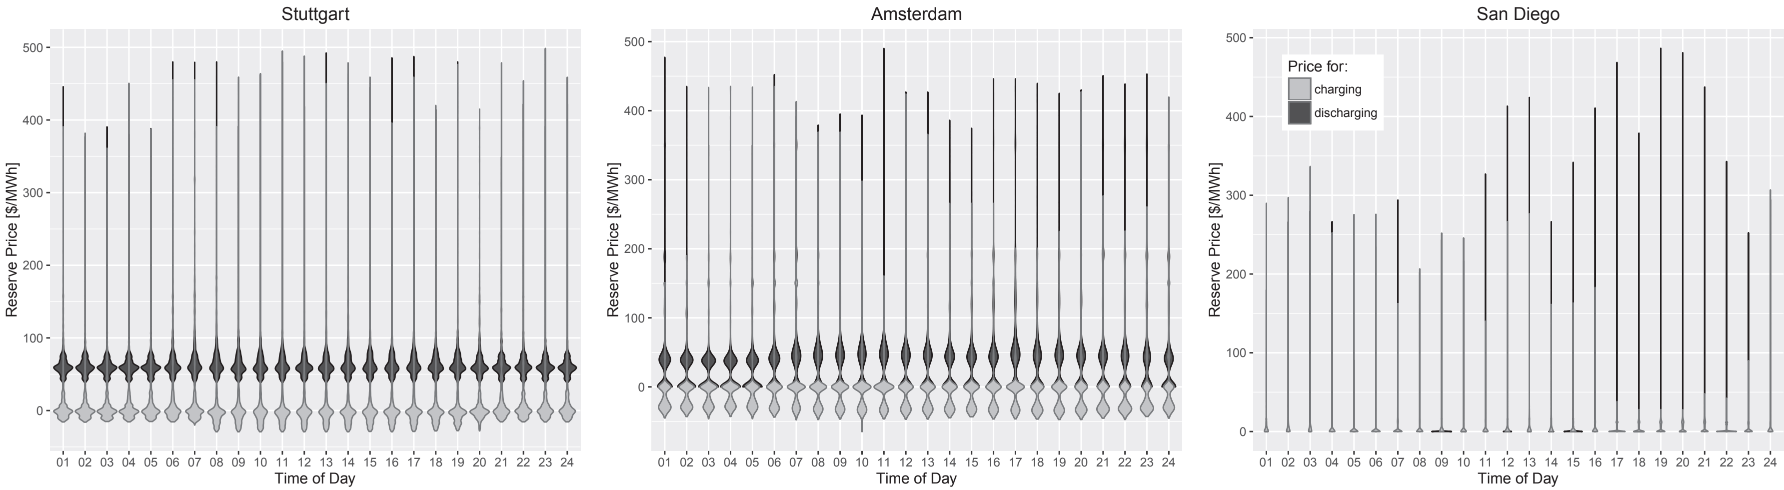

Figure 11: Regulation prices (June 30, 2014 to June 29, 2015), with standard deviation illustrating the extreme price volatility. Prices range from -60$\frac{\$}{MWh}$, at which the fleet can charge, to +500$\frac{\$}{MWh}$, at which energy can be sold back to the grid.

in Amsterdam and Stuttgart. The additional benefits from using the fleet as a VPP (grey line) consistently increases profits during the entire year but is more pronounced in Stuttgart and San Diego. If we would break down the annual gross profits per EV at the example of Stuttgart, fleet owners make \$3,900 from the rentals of an average EV, save \$163 by charging at cheaper time intervals, and make an additional \$14 from discharging electricity to the grid.

Our model and algorithms make decisions to maximize gross profits without knowing which rental transactions are going to happen in the future. Due to this uncertainty, the algorithm necessarily makes errors in its bidding process. Errors in committing EVs to a VPP when they could be rented out are the most significant in terms of their impact on profitability due to the asymmetric pay-offs. The error matrix (Table 5) shows that the model performs extremely well in predicting when it would be more profitable to use EVs for rental transactions, rather than committing them to the VPP; in all three cities, errors on this occur less than 1% of the time. With the stratified sampling we are able to achieve this high accuracy for rental transactions as when the model is not very certain, it will make EVs available to be rented. However, this comes at the expense of not committing EVs to the VPP in periods when it would have been profitable. This happens between 42% and 80% of the time, depending on the misclassification weights per city. Table 6 shows that the decisions made have resulted in profits in absolute terms in all three cities. In particular, the earnings from EVs that act as VPPs exceed by a long way the opportunity cost of lost rentals. In Stuttgart, for example, Car2Go earned \$118,000 on top of its \$2m dollar rental business, while the rentals lost due to VPP commitments only caused losses of \$31,000. While 2,173 customer transactions could not be carried out immediately, the VPPs increased Car2Go's annual gross profits by 4.6% under the conditions observed in Stuttgart. Also in the other cities the use of FleetPower increased gross profits 1.7% in Amsterdam and 3% in San Diego. We attribute these differences as being due to the price levels in the case of Amsterdam (see Figure 11), and in San Diego as stemming from a lack of suitable infrastructure. We will analyze the effect of infrastructure in more detail in the next section. In all three cities the discharging capacity offered to the market was very seldom taken up by buyers (only 61 MWh was actually purchased for use), and in Amsterdam none of these asks were ever accepted. This was due to our high ask prices, which included the costs of battery depreciation. Currently 93% of the gross profits are earned by savings on the electricity bill for the EVs; in many cases the grid needs to urgently

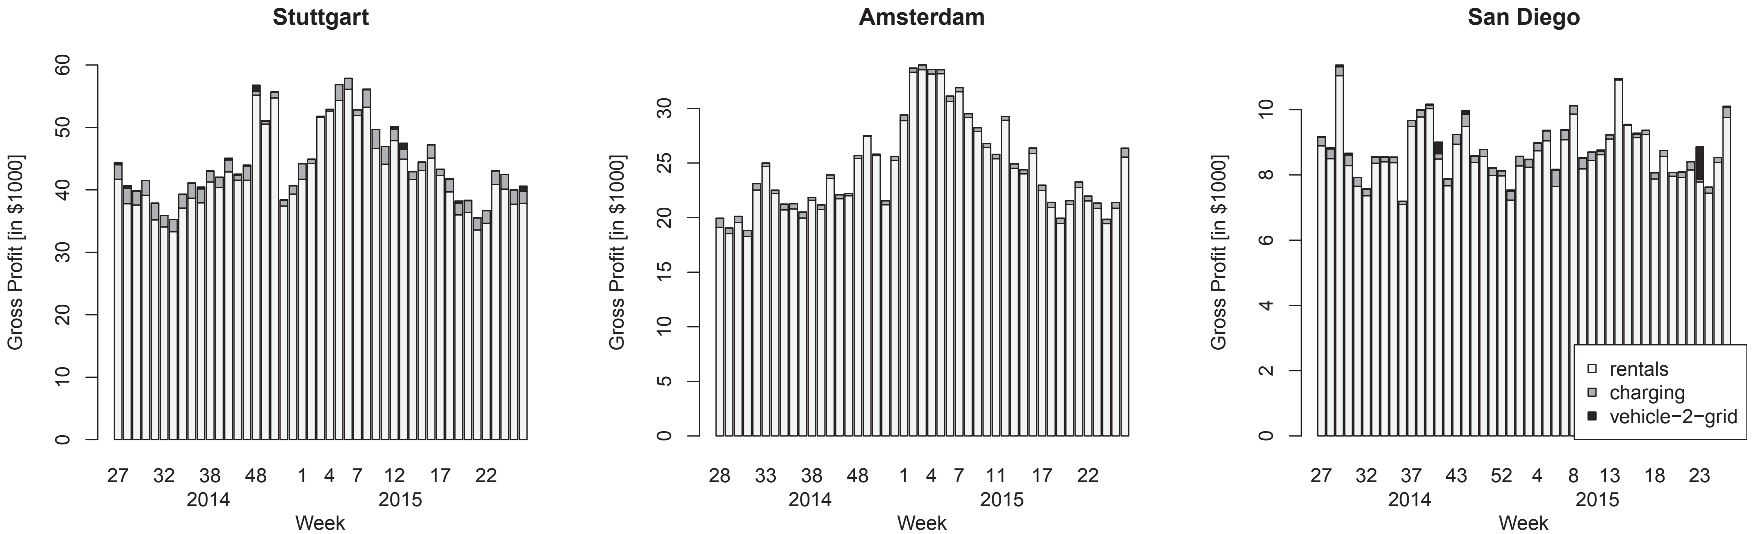

Figure 12: Shows the profits over the year. The profits from offering VPP power on the real-time market increases the gross profits of Car2Go consistently.

get rid of surplus electricity meaning that Car2Go actually gets paid to charge its EVs. For more information about the decision outcomes, see Table 7.

Impact on the Grid

Creating a VPP using an EV fleet provides a sound business case for a carsharing fleet. However, it is also beneficial to the grid and thereby society, because it provides additional reserve power to help keep the grid in balance at all times. This is already beneficial for the operation of the grid, but it becomes essential when a large proportion of weather-dependent renewable energy sources come on to the market. The VPP supports the grid by providing and consuming electricity on demand within seconds. The capacity that Car2Go provided to the market in each city is displayed in Figure 13. While Car2Go consumes a substantial amount of surplus energy (a negative value on the y-axis means the total quantity charged), it discharges its EVs only infrequently in Stuttgart and San Diego and never in Amsterdam. Due to the low discharging prices, the cost of battery wear cannot be covered, with the consequence that asks to discharge EVs are accepted infrequently. We will elaborate on the impact of larger fleets on the balancing market in Section .

Sensitivity Analysis: Nonrecurring Events

In weeks when there is a high volume of rental transactions the fleet earns less from VPP operation, and conversely it earns more from VPP when the rental business is not going well. This can be inferred from the downward sloping trend in all three cities in Figure 14. When there are fewer rental transactions the probability of EVs being available for VPPs is higher and the chance of forgoing rental revenues is lower. In consequence the FleetPower business model is a natural complement to the existing carsharing business model; it allows fleets to bridge periods when rentals are declining with VPP profits of up to 8%, 4%, and 13% in Stuttgart, Amsterdam, and San Diego respectively (see Figure 14). When analyzing the data we found several outliers when events occur that cannot be inferred from the training period. For example, there were European, regional, and council elections held in Stuttgart on Sunday, May 25, 2014. Polling stations were open from 8 am to 6 pm, and this led more people than usual to rent a car that Sunday morning between 8 am and 11 am to drive to the polling stations. Our algorithm was not aware of the elections and made commitments to the energy market as if for a normal Sunday. The result was that on that particular day the number of people who would have rented a car was exceptionally high, but we had

Table 5: Confusion (error) matrix (%): The matrix shows the accuracy of FleetPower's decisions. Because of the asymmetric pay-off the algorithm is trained not to bid for a VPP when rental transactions occur.

	Predicted					
Actual	**VPP**	**Rented**	**VPP**	**Rented**	**VPP**	**Rented**
VPP	53.00%	42.28%	28.11 %	68.12%	16.89%	80.27%
Rented	0.07%	4.64%	0.00%	3.76%	0.01%	2.83%
	[Stuttgart]		[Amsterdam]		[San Diego]	

Table 6: Confusion (error) matrix (in 1000 $): The matrix shows the monetary impact of FleetPower's decisions. The added value from VPP exceeds by far the losses from lost rentals in all cities.

	Predicted					
Actual	**VPP**	**Rented**	**VPP**	**Rented**	**VPP**	**Rented**
VPP	118	-	21	-	13	-
Rented	-31	1915	-1.4	1092	-0.7	396
	[Stuttgart]		[Amsterdam]		[San Diego]	

Table 7: Decision outcome results over a one-year period.

Description	**Stuttgart**	**Amsterdam**	**San Diego**
Discharged (V2G)			
Discharged energy sold (MWh)	54	0	7
Number of lost rentals	106	0	9
Increase in Gross profit (%)	0.4	0	0.6
Increase in Gross profit (in 1000 $)	7	0	2
Charged			
Quantity of energy bought (MWh)	1248	493	136
Number of lost rentals	2067	94	50
Increase in Gross profit (%)	4.2	1.7	2.4
Increase in Gross profit (in 1000 $)	80	19	9.5

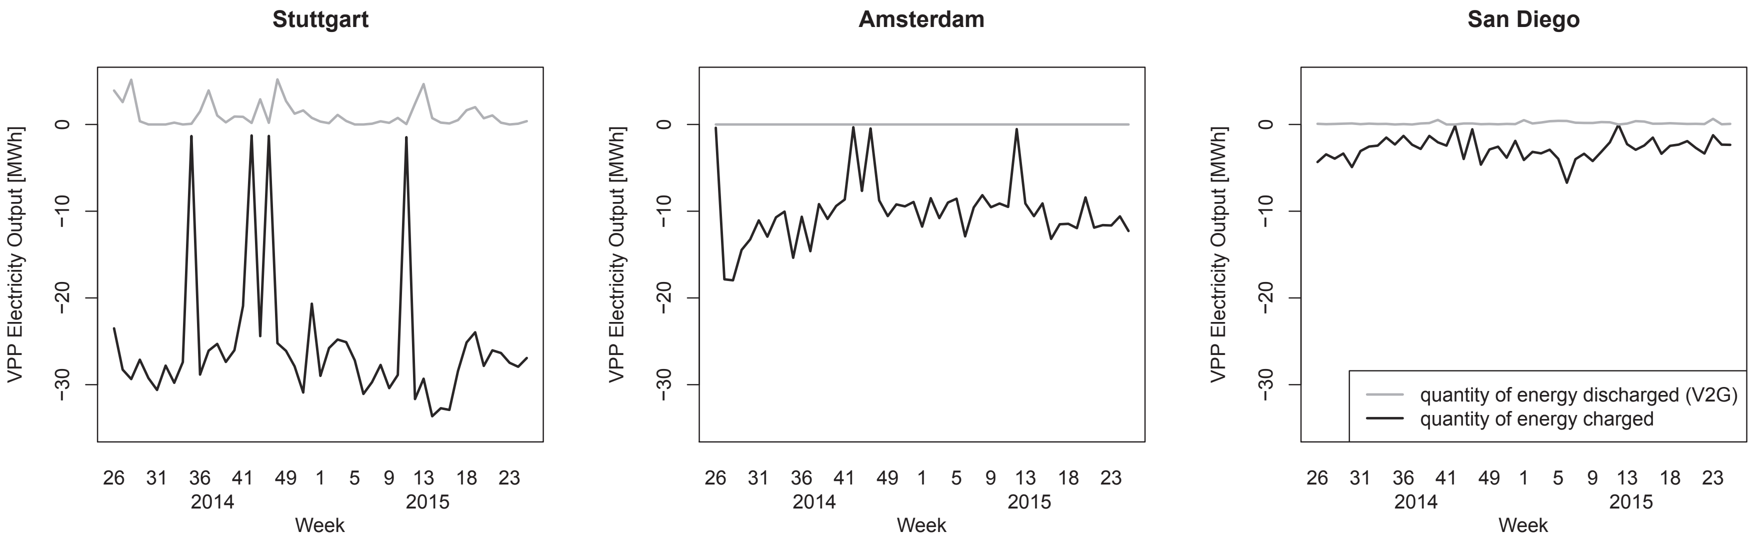

Figure 13: Shows the VPP output over the year. It is striking that vehicle-2-grid does not occur at all in Amsterdam, due to the low prices paid for discharging.

to turn 37 people away, with an opportunity cost of $513. A natural extension of our algorithm would be to incorporate known events and holidays in order to anticipate such events and increase the accuracy and profitability of our model. This concept of adjustable autonomy is covered in previous literature (Bichler et al., 2010); we therefore did not include it in our algorithm, which makes the profit estimations of our algorithms more conservative.

Sensitivity Analysis: Charging Infrastructure

Charging infrastructure plays an important role for the business model, as more charging stations suggest there will be a higher average state of charge, leaving more room for both the rental and VPP business. Because the charging infrastructure is still under development we artificially increase the number of charging stations by 50% (factor of 1.5) and 100% (factor of 2) to assess the impact of future developments on the VPP business model. To do this we add charging stations to the parking spots that are most popular in terms of the number of hours parked. We stop placing charging stations when the total number of hours that EVs are parked at a charging station is twice what it is now. Surprisingly, of the 100 most popular parking spots in each city, 91 in Stuttgart already had charging stations, whereas only 57 did in Amsterdam, and 38 in San Diego. This result is especially surprising with regard to Amsterdam, which has almost seven times as many charging stations as Stuttgart. This means that both Amsterdam and San Diego could significantly improve the quality of their infrastructure by putting in charging stations in a few key locations. Charging stations in the Danil Goedkoopstraat 14 in Amsterdam, the 342-398 Market Street in San Diego, and the Eichenwaldallee 202 in Stuttgart would yield the largest improvements in infrastructure for these cities. A 59% increase in the number of hours that EVs are parked at charging stations (achieved through strategic placement of additional stations) has a linearly positive effect on the gross profits of Car2Go. If this were done, the gross profits would be increased by 8%, 3%, and 5% for Stuttgart, Amsterdam, and San Diego respectively. If the number of charging stations was increased by 100%, the profits for each city would increase by 9%, 4%, and 6% respectively (see Table 8). These results indicate that the charging stations in Stuttgart are already in particularly good locations, so that additional charging stations do not enhance profits as much. However, increasing the number of charging stations would significantly enhance profits in San Diego especially, but also in Amsterdam.

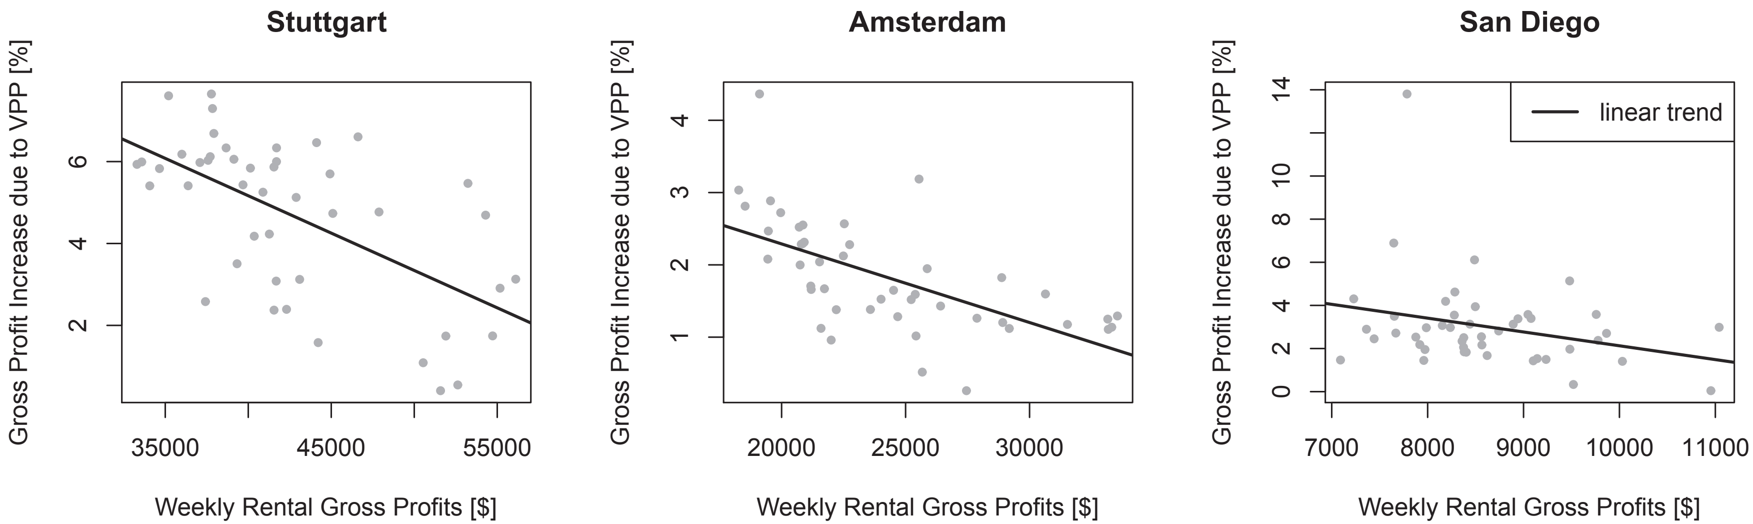

Figure 14: Shows the VPP profits as a function of the exogenous variable weekly rental profits. The downward sloping trend indicates that the profits from the VPPs complement periods where there are few rental transactions.

Table 8: Sensitivity analysis for charging station roll-out. Shows how the gross profit of Car2Go's virtual power plant business model increases as a function of charging stations with the growth factors 1 (as is), 1.5 (50% increase) and 2 (100% increase).

	Growth factor: 1	**Growth factor: 1.5**	**Growth factor: 2**
Stuttgart	↑ 4.6% gross profit	↑ 7.7% gross profit	↑ 8.7% gross profit
Amsterdam	↑ 1.7% gross profit	↑ 3.2% gross profit	↑ 3.8% gross profit
San Diego	↑ 3.0% gross profit	↑ 5.3% gross profit	↑ 5.5% gross profit

Sensitivity Analysis: Price Changes

As more renewable energy sources are brought into the energy mix in each of our three locations, more control reserve power (increase in $Q*$) will be needed to cover periods when these sources are either over- or underproducing. This increases the probability that higher priced bids from control reserves are also accepted - from discharging (V2G), for example. Therefore, we will analyse a scenario in which the need for secondary control reserves increases from the current demand Q^* to the future demand including more renewable energy sources $Q^{*'}$ (for both up- and down-regulation). We will model the future price change as illustrated for discharging on July 1 from midnight to 00:15 am in Figure 15. The increased demand causes the maximum acceptable price to rise from the current price P to the future price $P^{'}$. We can assess the difference in price discharge EVs between now and the future (ΔP) from the supply curves. The same happens to the demand for dumping surplus capacity; an increase in Q^* also changes the price to charge EVs. The price changes for both discharging and charging has a direct effect on the profitability of a VPP.

In Stuttgart we know the functions of the supply and demand curves, though this information was not available for San Diego and Amsterdam. We have therefore analyzed what effect changes in price have on profitability for Stuttgart. We consider a scenario in which the control reserve requirement Q^* doubles to $Q^{*'}$, and the impact of this is that the gross profits increase from 4.6% to 14.2%. This increase also has a positive effect on V2G usage; about two-thirds of the profits from VPPs are earned from V2G. The price increase makes discharging from EVs more cost-effective for fleets, as it means they are then reimbursed for the cost of battery depreciation.

Sensitivity Analysis: Scalability

The impact of a EV fleet of 350-500 cars on a country's real-time power reserve needs is negligible. However, a growth in the adoption of EVs raises the question of how this will impact the grid. Specifically we have analyzed what effect the number of

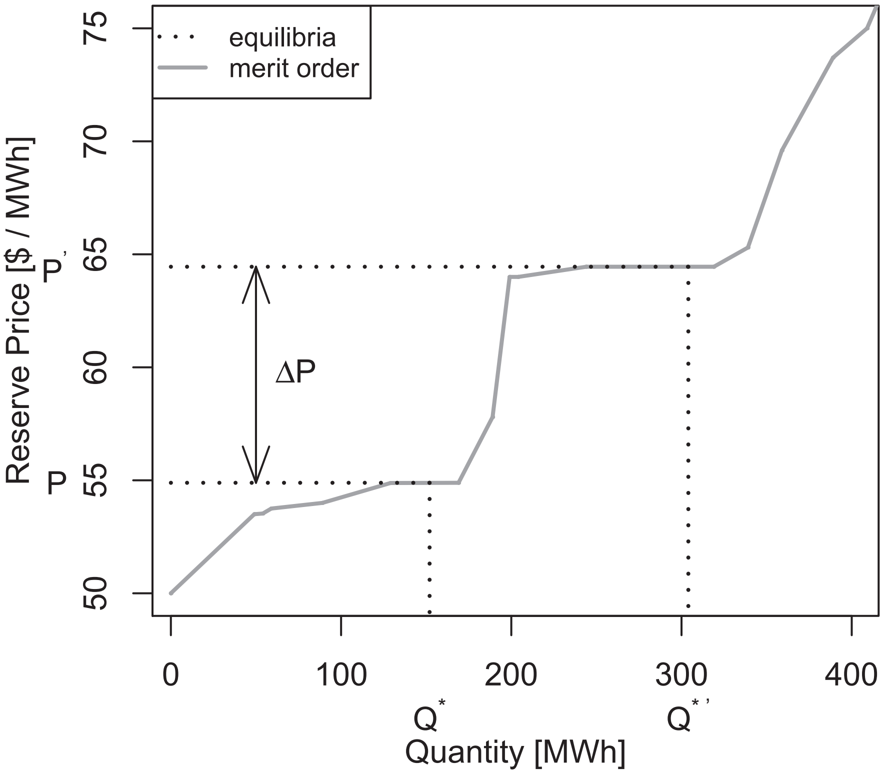

Figure 15: An increase in the reserve power requirement (from Q^* to $Q^{*'}$) shifts prices along the merit order (supply curve) to higher prices (from P to $P^{'}$) with a price difference of ΔP.

EVs might have in terms of their ability to provide sufficient real-time power reserves. For that reason we have modeled the frequency distributions of the capacity provided by the fleet during the four different seasons and time blocks of the day: {night time (midnight to 7 am), morning (7 am to 10 am), noon (11 am to 2 pm), afternoon (3 pm to 6 pm), and evening (7 pm to midnight)}. In all 20 time-block/season combinations, it is frequently the case that the fleet did not provide any reserve capacity at all. To model that, we first proportionally sampled whether a fleet provided no control reserve capacity at all. If it provided a capacity greater than 0, we modeled distributions for each time-block/season combination (all normally distributed with the Kolmogorow-Smirnow test at a 5% significance level) and then picked random samples from this distribution to determine the reserve power that could potentially be provided by a larger fleet or more fleets. However, for discharging in particular, large fleets currently also provide only a small share of the total power requirement of the real-time energy market. Even though the total storage capacity would be large enough, the priority given to renting out vehicles rather than charging and discharging at the current price levels would not allow the entire market to be balanced by EVs alone. However, if the demand for reserve power were to double, prices would reach a level that would justify using V2G on a large scale. We find that if half of Germany's cars were EVs, they would be able to deliver all the required reserve power needed to meet the demand during our test period from June 2014 to June 2015.

Conclusions

We have proposed and evaluated the FleetPower decision support system, which enables EV fleets to participate in the energy market as well as to continue their traditional rental business. We do this by using an intelligent agent that decides whether an EV at a specific location should be made available for rent, or whether it should be charged or discharged in form of a virtual power plant, providing an ancillary service. The system makes this decision based on forecasted rental transactions, charging, and discharging. Our tests show that using EVs for ancillary services consistently enhances gross profits for the EV fleet by 1.7% and 4.6%, depending on the location, representing an increase in annual gross profit of up to $86,000. This compares well to other studies, which found 14% in a stationary storage setting (Vytelingum et al., 2011), and up to 5% for electric vehicles (Schill, 2011). However, neither of these studies take into account that vehicles driving patterns are unknown a priori. In this study we have taken this probability of not serving rental customers into account with an asymmetric pay-off. We comparing the agent's performance across Stuttgart, Amsterdam, and San Diego, and are able to show how profitability is affected by the charging infrastructure in place and by energy prices on the regulation markets. V2G currently accounts for only a small proportion of these additional profits, as 90% of the profits come from electricity savings. However, we show that V2G has a strong impact on the gross profits of carsharing fleets when the demand for reserve power increases. Additionally, we demonstrate that the roll-out of additional charging stations in the future will have a positive effect on the business model, and we make recommendations on how GPS data on parking duration could be used to position these stations strategically. With this decision support system it is possible to replace carbon-intensive back-up capacity with clean energy storage, but as there are not yet enough EVs on the street, they need to be combined with other fast-response technologies such as biogas or hydropower in order to balance volatile renewable energy sources such as wind or solar.

References

A. e. a. Agricola. DENA ancillary services study 2030. Security and reliability of a power supply with a high percentage of renewable energy. Technical report, German Energy Agency, Berlin, Germany, 2014.

B. Avci, K. Girotra, and S. Netessine. Electric vehicles with a battery switching station: Adoption and environmental impact. *Management Science*, 61(4):772–794, 2015.

R. Bapna, P. Goes, and A. Gupta. Pricing and allocation for quality-differentiated online services. *Management Science*, 51(7):1141–1150, 2005.

R. Berk, L. Sherman, G. Barnes, E. Kurtz, and L. Ahlman. Forecasting murder within a population of probationers and parolees: a high stakes application of statistical learning. *Journal of the Royal Statistical Society*, 172(1):191–211, 2009.

M. Bichler, A. Gupta, and W. Ketter. Designing smart markets. *Information Systems Research*, 21(4):688–699, 2010.

T. Brandt, S. Wagner, and D. Neumann. Evaluating a business model for vehicle-grid integration: Evidence from germany. *Transportation Research Part D: Transport and Environment*, 50:488–504, 2017.

J. Collins, W. Ketter, and N. Sadeh. Pushing the limits of rational agents: the trading agent competition for supply chain management. *AI Magazine*, 31(2):63, 2010.

J. Firnkorn and M. Müller. What will be the environmental effects of new free-floating car-sharing systems? The case of car2go in Ulm. *Ecological Economics*, 70(8): 1519–1528, 2011.

J. Gallien and L. Wein. A smart market for industrial procurement with capacity constraints. *Management Science*, 51(1):76–91, 2005.

F. Geels. From sectoral systems of innovation to socio-technical systems: Insights about dynamics and change from sociology and institutional theory. *Research Policy*, 33(6):897–920, 2004.

S. Gregor and A. Hevner. Positioning and presenting design science research for maximum impact. *MIS quarterly*, 37(2):337–355, 2013.

G. He, Q. Chen, C. Kang, P. Pinson, and Q. Xia. Optimal bidding strategy of battery storage in power markets considering performance-based regulation and battery cycle life. *IEEE Transactions on Smart Grid*, 7(5):2359–2367, 2016.

A. R. Hevner, S. T. March, J. Park, and S. Ram. Design Science in Information Systems Research. *MISQ*, 28(1):75–105, 2004.

International Grid Control Cooperation. Information on grid control cooperation and international development. *IGCC price model - transmission system operator market information*, pages 1–19, 2014.

M. Kahlen, W. Ketter, and J. van Dalen. Balancing with electric vehicles: A profitable business model. In *22nd European Conference on Information Systems*, pages 1–16, Tel Aviv, Israel, June 2014.

E. Kara, J. Macdonald, D. Black, M. Berges, G. Hug, and S. Kiliccote. Estimating the benefits of electric vehicle smart charging at non-residential locations: A data-driven approach. *Applied Energy*, 155:515–525, 2015.

J. G. Kassakian and R. Schmalensee. The future of the electric grid: An interdisciplinary MIT study. Technical report, Massachusetts Institute of Technology, 2011. ISBN 978-0-9828008-6-7.

W. Ketter, J. Collins, and P. Reddy. Power TAC: A competitive economic simulation of the smart grid. *Energy Economics*, 39(0):262–270, 2013.

W. Ketter, M. Peters, J. Collins, and A. Gupta. A multiagent competitive gaming platform to address societal challenges. *MISQ*, 40(2):447–460, 2016a.

W. Ketter, M. Peters, J. Collins, and A. Gupta. Competitive benchmarking: An IS research approach to address wicked problems with big data and analytics. *MISQ*, 40(4):1057–1080, 2016b.

E. Kim, R. Tabors, R. Stoddard, and T. Allmendinger. Carbitrage: Utility integration of electric vehicles and the smart grid. *The Electricity Journal*, 25(2):16–23, 2012.

A. Knaut, F. Obermüller, and F. Weiser. Tender frequency and market concentration in balancing power markets. Technical Report 4, Energiewirtschaftliches Institut an der Universitaet zu Koeln (EWI), 2017.

C. Loock, T. Staake, and F. Thiesse. Motivating energy-efficient behavior with green IS: an investigation of goal setting and the role of defaults. *MISQ*, 37(4):1313–1332, 2013.

H. Mak, Y. Rong, and Z. M. Shen. Infrastructure planning for electric vehicles with battery swapping. *Management Science*, 59(7):1557–1575, 2013.

A. Malhotra, N. Melville, and R. Watson. Spurring impactful research on information systems for environmental sustainability. *MISQ*, 37(4):1265–1274, 2013.

E. Mashhour and S. Moghaddas-Tafreshi. Bidding strategy of virtual power plant for participating in energy and spinning reserve marketspart ii: Numerical analysis. *IEEE Transactions on Power Systems*, 26(2):957–964, 2011.

K. McCabe, S. Rassenti, and V. Smith. Smart computer-assisted markets. *Science*, 254(5031):534–538, 1991.

Nest Labs. Energy savings from the Nest learning thermostat: Energy bill analysis results. White paper, Palo Alto, California, United States, 2015.

M. E. Porter and J. E. Heppelmann. How smart, connected products are transforming competition. *Harvard Business Review*, 92:11–64, 2014.

D. Pudjianto, C. Ramsay, and G. Strbac. Virtual power plant and system integration of distributed energy resources. *Renewable Power Generation, IET*, 1(1):10–16, 2007.

S. Reichert. Considerations for highly efficient bidirectional battery chargers for e-mobility. In *E-Mobility. Technologien, Infrastruktur, Maerkte*, page 5, Leipzig, 2010. VDE-Verlag. ISBN 978-3-8007-3304-0.

C. Robusto. The cosine-haversine formula. *American Mathematical Monthly*, 64(1): 38–40, 1957.

W. Schill. Electric vehicles in imperfect electricity markets: The case of Germany. *Energy Policy*, 39(10):6178–6189, 2011.

R. Sioshansi. The impacts of electricity tariffs on plug-in hybrid electric vehicle charging, costs, and emissions. *Operations Research*, 60(3):5060–516, 2012.

J. Tomic and W. Kempton. Using fleets of electric-drive vehicles for grid support. *Journal of Power Sources*, 168(2):459–468, 2007.

K. Valogianni, W. Ketter, J. Collins, and D. Zhdanov. Enabling sustainable smart homes: An intelligent agent approach. *35th International Conference on Information Systems (ICIS)*, pages 1–20, 2014.

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

S. Wagner, T. Brandt, and D. Neumann. In free float: Developing business analytics support for carsharing providers. *Omega*, 59(0):4–14, 2016.

R. T. Watson, M. Boudreau, and A. J. Chen. Information systems and environmentally sustainable development: energy informatics and new directions for the IS community. *MISQ*, 34(1):4, 2010.

A. Wolfson, D. Tavor, S. Mark, M. Schermann, and H. Krcmar. Better place: a case study of the reciprocal relations between sustainability and service. *Service Science*, 3(2):172–181, 2011.

Y. Zhou, A. Scheller-Wolf, N. Secomandi, and S. Smith. Electricity trading and negative prices: Storage vs. disposal. *Management Science*, 0(0):1–19, 2015.