

Theoretical and Practical Foundations of Large-Scale Agent-Based Micro-Storage in the Smart Grid

Perukrishnen Vytelingum
Thomas D. Voice
Sarvapali D. Ramchurn
Alex Rogers
Nicholas R. Jennings

Agents, Interaction and Complexity Group
School of Electronics and Computer Science
University of Southampton
Southampton, SO17 1BJ, UK.

PV@ECS.SOTON.AC.UK
TDV@ECS.SOTON.AC.UK
SDR@ECS.SOTON.AC.UK
ACR@ECS.SOTON.AC.UK
NRJ@ECS.SOTON.AC.UK

Abstract

In this paper, we present a novel decentralised management technique that allows electricity micro-storage devices, deployed within individual homes as part of a smart electricity grid, to converge to profitable and efficient behaviours. Specifically, we propose the use of software agents, residing on the users' smart meters, to automate and optimise the charging cycle of micro-storage devices in the home to minimise its costs, and we present a study of both the theoretical underpinnings and the implications of a practical solution, of using software agents for such micro-storage management. First, by formalising the strategic choice each agent makes in deciding when to charge its battery, we develop a game-theoretic framework within which we can analyse the competitive equilibria of an electricity grid populated by such agents and hence predict the best consumption profile for that population given their battery properties and individual load profiles. Our framework also allows us to compute theoretical bounds on the amount of storage that will be adopted by the population. Second, to analyse the practical implications of micro-storage deployments in the grid, we present a novel algorithm that each agent can use to optimise its battery storage profile in order to minimise its owner's costs. This algorithm uses a learning strategy that allows it to adapt as the price of electricity changes in real-time, and we show that the adoption of these strategies results in the system converging to the theoretical equilibria. Finally, we empirically evaluate the adoption of our micro-storage management technique within a complex setting, based on the UK electricity market, where agents may have widely varying load profiles, battery types, and learning rates. In this case, our approach yields savings of up to 14% in energy cost for an average consumer using a storage device with a capacity of less than 4.5 kWh and up to a 7% reduction in carbon emissions resulting from electricity generation (with only domestic consumers adopting micro-storage and, commercial and industrial consumers not changing their demand). Moreover, corroborating our theoretical bound, an equilibrium is shown to exist where no more than 48% of households would wish to own storage devices and where social welfare would also be improved (yielding overall annual savings of nearly £1.5B).

1. Introduction

The vision of an intelligent electricity delivery network, commonly called the *smart grid*, has been advocated as one of the main solutions to ensuring sustainable energy provision

(US Department Of Energy, 2003; Galvin & Yeager, 2008; DECC, 2009). Such grids aim to reduce inefficiencies in energy usage, minimise carbon emissions and reduce costs to generate electricity. A key element of this vision is that consumers should be able to respond to the current condition of the grid, by shifting or reducing their use of electricity, through a two-way communication channel between them and the other actors in the system (i.e., suppliers, other consumers, and grid operators). By so doing, it is expected that peaks in demand can be reduced, which, in turn, would reduce the need for expensive and carbon-intensive peaking plants (i.e., spinning reserve) that rely on fossil fuels.

Two of the key technical enablers of the smart grid are increased degrees of automation and the ability to store energy. In the former case, consumers need to be able to delegate the complex and time-consuming reasoning about shifting or reducing demand, subject to their individual preference, to software agents that will act on their behalf (Ramchurn, Vytelingum, Rogers, & Jennings, 2011c). To this end, *smart meters* have been developed to provide consumers and their agents with real-time information about a home's consumption and the state of the grid, and to provide suppliers and grid operators with consumption data from homes. Moreover, smart meter roll-outs have now been mandated in a number of countries, including France (by 2016), Spain (by 2018) and the UK (by 2020). Through such information feeds, consumers should be able to improve their management of energy (e.g., switching off devices they do not need or rescheduling power-hungry devices to other times). In the latter case, the agents can use the information to decide to store electricity at times when overall demand from the grid is low (and generally cheaper) and re-use this electricity when the grid is operating close to its limits (i.e., when generators are operating near capacity or transmission lines are close to overload and electricity is generally more expensive). Furthermore, storage devices can also be used to compensate for the variability of many forms of renewable electricity generation (e.g., wind, wave, solar) that is likely to be an increasingly prominent component of the future grid (Bathurst & Strbac, 2003).

To date, most research on storage technologies has focused on designing new low-cost large-scale storage devices (with capacities of the order of 10-100MWh) that are able to efficiently store electricity for long periods of time and allow a sufficient number of charging/discharging cycles without significant degradation in performance (Bathurst & Strbac, 2003). However, with the development of a large variety of micro-storage devices (i.e., with capacities of the order of kWh) that can be installed in homes¹ and vehicles², a future where large numbers of individual consumers can store small amounts of electricity in order to accommodate peaks in demand and variability in supply will soon be possible. There are, however, a number of potential challenges in this setting. First, if all consumers decide to charge their batteries *at the same time*, because the price is low, a significant peak in demand could well ensue. This would, in turn, result in higher electricity generation costs and greater carbon emissions and could overload a system that is already operating close to its maximum capacity (resulting in a brown-out or, in the worst case, a black-out). Indeed, such unintended population-wide synchronisation has already been seen in real-world

-
1. See batteries recently developed by GS Yuasa (<http://lithiumenergy.jp/en/products/>) or Power Yiile (<http://eliyypower.co.jp/english/lithium-ion/>).
 2. Vehicle to grid (V2G) technologies enable energy to be stored in the batteries of electric vehicles (EVs) (e.g., Mini-E or Nissan Leaf) or plug-in hybrid electric vehicles (PHEVs) (e.g., Toyota Plug-in Prius or Chevrolet Volt) (Sovacool & Hirsh, 2009; Lund & Kempton, 2008).

demand-response trials where consumers manually react to critical pricing periods (Hammerstrom et al., 2008). Second, if consumers were to simply charge their batteries to ensure they have sufficient energy to cover their *whole demand* across a day, they may end up paying more than they need to. This is because it may be cheaper to use the grid-supplied electricity at certain times, rather than the energy already stored in the battery. Finally, if most homes in the system start using storage devices and manage to reduce peak demand (by optimising their battery charging and usage costs), electricity prices may become lower than the price of stored electricity (including the cost of the battery), thus voiding the need for micro-storage and breaking down the market for such devices.

Addressing the aforementioned issues through the formulation of conventional closed-form solutions (see Section 2 for more details) is challenged by the fact that the system is composed of large numbers of distinct stakeholders (typically millions of consumers and tens of market makers and network managers) operating in a completely decentralised fashion where they individually act to satisfy their own particular objectives and constraints (to supply or use energy) which may conflict (e.g., the network managers aim to minimise peaks in demand while consumers aim to minimise their costs and use devices at the most convenient time of the day). Against this background, the agent-based approach can be used both as a framework to analyse the properties of such systems and also as an implementation technology (Exarchakos, Leach, & Exarchakos, 2009; Houwing, Negenborn, Heijnen, Schutter, & Hellendoorn, 2007; van Dam, Houwing, & Bouwmans, 2008; Rogers & Jennings, 2010). In particular, game theory can be used to determine the properties of the system as the multiple self-interested parties interact and software agents can be installed on the smart meters to optimise the usage and storage profile of the house using information from a variety of sources (e.g., weather data to predict heating needs and costs or price plan data from suppliers). Now, most of the existing approaches to applying intelligent agents study how individual homes could optimise the way they store energy or how storage devices could coordinate with renewable energy generation facilities to maximise energy used from such sources (see Section 2 for more details). However, in so doing, they ignore the individual selfish preferences of each consumer and do not model well the real impact of agents learning to adapt to the constraints that they themselves impose on the system. To remedy this, it is crucial to devise an approach that focuses on the system dynamics where all agents in the system are given the freedom to buy electricity whenever they see fit.

In this paper, we take just such an approach and provide both an analytical and a practical solution to the decentralised control of micro-storage devices in the grid. Specifically, we develop a game-theoretic framework for modelling storage devices in large-scale systems where each device is controlled by a self-interested agent that aims to maximise its monetary profits when a real-time price for electricity is provided by the grid. Using this framework, under certain reasonable assumptions, we are able to predict the equilibria of the system given that each agent behaves rationally (i.e., always adopts a storage profile that minimises its costs) and only reacts to a price signal. Given this, we go on to devise intelligent agent-based storage strategies that can learn the best storage profile given market prices that are themselves a result of the aggregate storage and consumption profile of all the agents in the system. Crucially, we show that agents using such strategies achieve our predicted equilibria and, building upon this, we simulate large populations of agents in order to predict the system behaviour under various conditions.

In somewhat more detail, this work advances the state of the art in the following ways:

1. We provide a novel analytical game-theoretic framework that captures the synchronous behaviours of consumers with micro-storage devices within the smart grid. We use this to study electricity storage strategies that agents might adopt in a smart grid with real-time pricing. Given the typical electricity usage profiles of consumers, we are able to compute the competitive equilibria which describe when each individual agent is going to charge its device, use its stored electricity, or use electricity from the grid. We then use this analytical solution to benchmark our decentralised solution.
2. We provide a theoretical bound on the storage capacity that will be needed by the population, as well as a bound on the portion of the population that will adopt storage.
3. We provide new micro-storage strategies that enable agents to learn the best storage profile to adopt, even taking into account the probable heterogeneity of the other micro-storage devices adopted by consumers across the grid (e.g., these devices will vary in capacity and in how fast they can be charged or discharged). Our practical strategies are shown to converge to the same competitive equilibria as those predicted by our analytical framework and come with system-wide benefits that include reduced carbon emissions, as well as cost savings.
4. We show how agents, having learnt their best storage profile, can also learn to buy the most profitable storage capacity. Given this, using evolutionary game theoretic analysis, we are able to predict the portion of the population that would actually acquire such storage capacity to maximise their savings. We show that this is not the entire population. Rather, it is just under half of them, confirming the theoretical bounds that our analytical framework predicts.

In short, this is the first attempt at modelling, predicting equilibria and building intelligent strategies for the problem of electricity storage on a large scale. Our approach also justifies and provides the basis for the implementation of real-time electricity pricing for domestic electricity distribution.

The rest of this paper is structured as follows. In Section 2, we discuss related work in the area of electricity storage and electricity markets. Section 3 discusses the key features of such markets and lays down the general assumptions upon which we build our framework. Section 4 presents our game-theoretic framework and shows how the competitive equilibria of the system can be computed and, based on such equilibria, determines how many users are likely to adopt micro-storage. Building on this, Section 5 describes the dynamics of a market where agents are given the ability to learn their best storage profile and Section 6 empirically studies this system through simulations. Then, Section 7 expands on the cost-benefit analysis of storage in a system given a heterogenous population of consumers with different usage profiles. Finally, Section 8 concludes.

2. Background

In this section, we first review current storage technologies to illustrate how micro-storage has evolved to date and is likely to be a potent energy management technology in the

future. We follow with a discussion of existing approaches in the power systems literature to using storage, including agent-based approaches that have been proposed to manage the grid in general and micro-storage in particular. Furthermore, we motivate the need to have real-time pricing mechanisms in the grid in order to enable more responsive demand by considering the response of consumers to fixed prices and time-of-use tariffs.

2.1 Energy Storage Technologies

The large scale storage of electricity within an electricity grid is not a new concept. Many countries operate pumped-storage power plants which store electricity by pumping water to a high reservoir when prices are low, and then release this water and use it to generate electricity when prices are high or supply is short. For example, Dinorwig Power Station in the UK can store approximately 10GWh of electricity, and can supply this stored energy over six hours, while providing up to 1.8GW of power (Williams, 1984).

However, the increased need for storage capacity as electricity grids increasingly seek to incorporate intermittent renewable sources, such as wind generation or solar power, coupled with the high capital costs of such pumped storage plants and the limited number of suitable sites for their construction, mean that much recent research has focused on alternative smaller storage solutions which typically store 10-100MWh of electricity. Examples of technologies that have already been demonstrated commercially include the use of underground caverns to store electricity by compressing and releasing air (Swider, 2007) and its storage in large chemical flow batteries (Shibata & Sato, 1999).

Most recently, attention has turned to micro-storage³ of up to 20kWh of electricity which might be installed within homes as part of a smart grid roll-out. This interest has been largely driven by the rapidly decreasing cost of efficient batteries, such as lithium-ion cells or the nickel-zinc alternative, as they are developed for use within electric vehicles (den Bossche, Vergels, Mierlo, Matheys, & Autenboer, 2006). Indeed, the energy storage capacity required for a viable electric vehicle is close to the daily consumption of a home (e.g., the Chevrolet Volt battery has a capacity of 16 kWh and the Nissan Leaf can store up to 24 kWh). Thus, it is now possible to envisage that micro-storage devices will be widely used in the short to medium term, either as dedicated home storage batteries, or as an additional capability of electric vehicles.

2.2 Agent-Based Systems for the Grid

The use of software agents, residing on smart meters, was first envisaged by Schweppe, Tabors, Kirtley, Outhred, Pickel, and Cox (1980) who proposed mechanisms for agents to manage the use of electricity within a single home and to buy electricity in real-time on behalf of their owner. Since then, a number of agent-based approaches have been developed where multiple autonomous agents represent the interests of different actors on the grid. For example, Ygge, Akkermans, Andersson, Krejic, and Boertjes (1999) initiated work on abstracting electricity markets as multi-commodity markets and showed how agents trad-

3. Note that the average cost of a micro-storage battery today varies from around £300 to up to £1000 per kWh with an average start-up cost of £200. Note that these prices are only indicative and given the innovative nature of the area and the investment in electric vehicle battery technology, the cost of micro-storage is gradually decreasing (see <http://www.pikeresearch.com>).

ing energy for different times of the day, could generate efficient allocations . Moreover, Jennings, Corera, Laresgoiti, Mamdani, Perriolat, Skarek, and Varga (1996) developed coordination mechanisms for different actors on the grid to manage the allocation of transmission capacity and, more recently, Sun and Tesfatsion (2007) and Li and Tesfatsion (2009) developed agent-based electricity market simulations which incorporate transmission constraints and different types of buyers and sellers in the grid. Kok and Venekamp (2010) take a similar approach and implement an agent-based architecture to run the electricity market where the individual actors represent either generators or devices in the home. Note, however, that none of these approaches consider the daily consumption profile of consumers, nor how an agent might optimise its consumption and storage of electricity to maximise its owner's benefit. We believe this is crucial because agents have to be profitable to their users, each with her own specific needs and lifestyle, in order to be commercially viable. Specifically, agents should be able to provide personalised home energy management if they are to be deployed in the real world.

In this context, we note the early work of Daryanian, Bohn, and Tabors (1989) that illustrated how individual agents could optimise, through iterative algorithms, the load profile of a house using an electricity storage device. Their approach was, however, limited to considering very basic battery properties and did not consider wider issues for the grid such as the level of adoption of batteries in the population and the optimum storage capacity required for maximum savings. More recently, Houwing et al. (2007) provided algorithms for agents to optimise storage using small domestic combined heat and power (CHP) plants, but they ignored how populations of such agents would impact on the grid. On the other hand, Exarchakos et al. (2009) and van Dam et al. (2008) have studied the application of storage devices on a wider scale. They showed that using demand-side management techniques, where the storage profile of a number of homes is controlled centrally, can increase savings made in the system. Unfortunately, such a centralised approach automatically introduces a single point of failure in the system and does not take into account the individual preferences of each user to buy, use, or turn off her storage device. On the other hand, Ramchurn, Vytelingum, Rogers, and Jennings (2011b) consider a supplier retailing to a large number of agents that continuously optimise their storage, but they assume that the effect of such large-scale optimisation does not influence the wholesale price of the market which is usually not the case when considering a large enough proportion of the population. Thus, in this paper, we take a market-based approach similar to Ygge et al. (1999) and Ygge and Akkermans (1999), that is informationally decentralised (i.e., no centre with complete and perfect information is required), and therefore more robust, where each agent buys electricity in real-time markets and individually controls the storage profile of its associated home, based on real-time prices that reflect the demand (and supply) of the market.

2.3 Electricity Pricing Mechanisms

Most of the above approaches assume the existence of a control signal that dictates when the home must store and when it is best to use the stored electricity, thereby reducing the load on the grid at peak times (see Figure 1 for an example of the average UK demand that result in peaks) and making a saving when such reductions come with monetary rewards. To this end, Schweppe et al. (1980) proposed the use of real-time pricing (RTP)

or spot pricing of electricity as a better way to manage demand as conditions on the grid are accurately reflected by the price (as set by a single utility or an electricity market). Thus, contrary to the pricing signals that are traditionally used in the grid, that are either based on a fixed price or a time-of-use (TOU) price, whereby a premium is charged during times of anticipated peak demand, a real-time price reflects the current and continuously changing balance of supply capacity and demand and, even in some cases, the congestion on the network (e.g., locational marginal pricing approaches to pricing electricity at different points in the grid – see Harris, 2005). Unfortunately, Schweppe et al.’s solution was never fully implemented on a large scale at the time it was proposed. This was for a number of reasons. First, the properties of their solution were mainly proven analytically, under the general assumption that most agents will behave in a similar fashion and did not attempt to model the strategic choices that agents may make in charging their batteries (e.g., always charging at times they predict will be cheaper and always using their battery when they predict prices will be higher or charging their battery early for a whole day’s consumption to avoid price peaks later). Instead, in this paper, we develop a game-theoretic framework that fully captures the agent’s strategic behaviour within the context of the smart grid and we complement this approach with simulations in order to evaluate the performance of the system with a highly heterogeneous population of agents. Second, Schweppe et al.’s design also came up against problems associated with high communication costs and a lack of autonomous storage management technology. However, recent advances in computational power that make the deployment of autonomous agents entirely feasible, and new information and communication technologies such as wireless broadband internet and home energy management systems (e.g., AlertMe⁴ or Intel’s Home energy dashboard⁵) mean that real-time pricing for the domestic sector looks closer to being realised. Moreover, the financial commitment of countries such as the UK (£8.6M invested in smart metering infrastructure – see DECC, 2009) and the US (with 57.9 million smart meters planned for installation⁶) to the implementation of smart metering infrastructure, provides unprecedented support for the implementation of RTP.

Real-world trials, such as those of the GridWise alliance in the US (Hammerstrom et al., 2008) or the Energy Demand Research Project in the UK (Smith, 2010), and theoretical work such as those by Ramchurn, Vytelingum, Rogers, and Jennings (2011a) show that the more accurate RTP signals (i.e., representing real costs as opposed to the TOU pricing scheme) allow consumers to reduce their peak demand (and the duration of such peaks) by reacting more frequently to a more accurate 30-min-tariff pricing model (rather than over the peak and off-peak prices of TOU). By reducing such peaks, consumers under RTP can make significant savings compared to those under the TOU pricing. However, it is generally the case that more (short-duration) peaks exist in the RTP than the Fixed-Pricing model (see Ramchurn et al., 2011, where two long-duration peaks exist in the morning and evening), requiring the usage of expensive peaking plants only for short periods. This is because if the agents predict (based on previous days) a low price for the next day at a

4. <http://www.alertme.com/>.

5. <http://edc.intel.com/Applications/Energy-Solutions/Home-Energy-Management/>.

6. <http://www.pikeresearch.com/newsroom/>

57-9-million-smart-meters-currently-planned-for-installation-in-the-united-states




Figure 1: Average load profile (weekday in summer) in UK over half-hourly periods based on a fixed pricing model.

certain time, they will all turn on their devices, which then results in a peak in demand at that time. When such a mechanism is rolled out on a large scale, such *reactive* behaviours can cause unpredictable and significant peaks (compared to two predictable ones for TOU) in demand and prices, which, in turn, result in higher costs for individuals and greater stress on grid resources (transmission lines reaching their thermal limit and generators reaching their capacity).

Thus, if not properly managed, storage systems can be unprofitable for consumers and adversely impact the whole system (Holland, 2009; Williams & Wright, 1991; Bathurst & Strbac, 2003). Hence, in the setting we consider, it is important to know whether micro-storage can be individually beneficial and what strategies maximise the consumer's savings. It is also important to understand the system-wide effects of such strategies; in particular, quantifying the limits on the usefulness of small scale storage from a grid efficiency point of view and determining how different types of storage (with different costs and efficiencies) will be integrated into the system. These are the key open questions that are addressed by this paper (see Sections 4 and 5 respectively).

3. Model Description

As already noted, in this paper we seek to analyse the behaviour and impact of micro-storage devices from a large scale multi-agent systems point of view. That is, we consider the situation where large numbers of autonomous agents each control energy storage for a home, and interact with electricity suppliers within a market, with the aim of minimising

costs for their individual owners. Such a setup allows us to make a number of modelling assumptions, which we will detail in this section. Our analysis considers fixed time intervals consisting of single days, each separated into $T = 48$ settlement periods of half an hour. Each day, agents consume electricity which is bought from suppliers through an electricity market. This market operates for each time interval in the day, so that variations in demand over time can be met. The agents autonomously control the charging and discharging behaviour of their storage device in order to maximise their user's profit. We describe this process more thoroughly in Section 3.1. In our analysis, the market itself is modelled abstractly, following a macro-model which we give details of in Section 3.2. In order to measure the impact of the behaviour of energy storage agents in a wider context, we employ a number of metrics for grid efficiency that are described and explained in Section 3.3. For a table of notation definitions used throughout this paper, see Table 1.

3.1 Agents

We consider a set of N consumers, \mathcal{A} , which we define as self-interested agents that always aim to minimise their individual costs. Each agent $a \in \mathcal{A}$ has a load profile⁷ $\ell_i^a \forall i \in \mathcal{I} = \{1, \dots, T\}$, such that ℓ_i^a is the amount of electricity required by agent a for time interval i during each day. The aggregate load profile of the system is given by $\sum_{a \in \mathcal{A}} \ell_i^a = \ell_i$. We consider this load profile to be fixed over different days (although in practice there are seasonal variations in demand, there is a high degree of consistency from day to day). Each agent $a \in \mathcal{A}$ may also have some storage available to it, with capacity κ^a , efficiency η^a and running costs μ^a . Here, the cost μ^a represents the ongoing storage costs (for example, some battery devices expend energy through heating while they are in use or lose efficiency through the depletion of chemicals used in them). We do not incorporate any fixed capital investment by a at this stage, as these costs are fixed, and only the running costs can have any effect on which storage profile is most profitable. However, such fixed capital investments are important when users decide whether or not to purchase a battery, and we include them in the cost-benefit analysis in Section 7. The storage efficiency η^a and running cost μ^a are modelled to be such that if c amount of energy is stored, then only $\eta^a c$ may be discharged and the storage cost is $\mu^a c$.

In order to minimise costs, agent a can attempt to strategise over its storage profile, $b_i^a \forall i \in \mathcal{I}$, where for all $i \in \mathcal{I}$, $b_i^a = c_i^a - d_i^a$, where $c_i^a = (b_i^a)^+$ is the charging profile and $d_i^a = (b_i^a)^-$, is the discharging profile. Here, and throughout the paper, we use the notation $(\cdot)^+$ to denote positive part, that is, $y = (x)^+$ means $y = x$ if $x > 0$, $y = 0$ otherwise. Likewise we use $(x)^-$ to denote $(-x)^+$. These definitions implicitly assume that a user will not both charge and discharge her storage device over a single time interval. However, in our model, prices are fixed over each time interval, so we discount such behaviour as it can never be profitable. For each agent a , a feasible storage profile $b_i^a \forall i \in \mathcal{I}$ must satisfy $-D^a \leq b_i^a \leq C^a$, where D^a is the maximum discharging rate of the storage device,

7. We assume that the consumer's load is inelastic and, thus, insensitive to price changes. In reality, we would expect that the consumer's load shows slight elasticity, i.e., the consumer will reduce her demand if price increases and likewise increase her demand if price decreases. However, because demand elasticity of domestic consumers is generally small, we believe that the results presented in this paper still provide a good guide to the behaviour of real markets.

Notation	Definition
N	Number of agents in the system
\mathcal{A}	Set of all agents in the system
T	Number of time intervals in a day
\mathcal{I}	Set of time intervals
b_i^a	Net charge/discharge from storage device of agent a during time period i
c_i^a	Amount charged by agent a during time period i
d_i^a	Amount discharged by agent a during time period i
l_i^a	Electricity used by agent a during time period i
q_i^a	Electricity purchased by agent a during time period i
η^a	Efficiency of storage device for agent a
κ^a	Capacity of storage device for agent a
μ^a	Storage cost for agent a
c_0^a	Energy agent a has stored at the start of the day
C^a	Maximum (per interval) charge for agent a
D^a	Maximum (per interval) discharge for agent a
O_i^a	Maximum usable discharge for agent a during i
b_i	Net charge/discharge by all agents during i
c_i	Amount charged by all agents during i
d_i	Amount discharged by all agents during i
l_i	Total consumer load by all agents during i
q_i	Total electricity purchased by all agents for i
κ	Total storage capacity of the population
η	Homogeneous storage device efficiency
μ	Homogeneous storage costs
C	Maximum total per interval charge
D	Maximum total per interval discharge
O_i	Maximum total useful discharge during i
s	Supply function for wholesale market
p_i	Price of electricity during i
p^+	Charging price point
p^-	Discharging price point
$\Phi_i(\cdot), \Psi_i(\cdot)$	Equilibrium price functions for time interval i (see Definition 1)
$(\cdot)^+$	Positive part
$(\cdot)^-$	Negative part
LU	Low-end energy users
HU	High-end energy users
r	Agent's strategy to adopt or not storage
π_r^U	Probability of an agent of type $U \in \{LU, HU\}$ adopting strategy r
$u_U(\cdot, \cdot, \cdot)$	Payoff of an agent of type $U \in \{LU, HU\}$

Table 1: Notation definitions.




Figure 2: Half-hourly UK wholesale real-time buy prices (in the balancing market) plotted against total demand for August and September 2009 illustrating the correlation between price and demand. The dotted line represents the average of the points.

and C^a , the maximum charging rate. Since we are attempting to model the effect of the widespread adoption of micro-storage devices, we can assume that ℓ_i^a , C^a , and D^a are small in comparison to ℓ_i (since we are considering an electricity grid consisting of a large number of consumers). We denote the total storage capacity as $\kappa = \sum_{a \in \mathcal{A}} \kappa^a$, and the net storage profile as $b_i \forall i \in \mathcal{I}$ where $b_i = \sum_{a \in \mathcal{A}} b_i^a$. For each $i \in \mathcal{I}$ and $a \in \mathcal{A}$ we use $q_i^a = b_i^a + \ell_i^a$, to denote the net quantity of electricity purchased by a during i , and we use $q_i = b_i + \ell_i$ to denote the net quantity of electricity purchased by all users during interval i . The aggregate maximum charging and discharging rates are defined as $C = \sum_{a \in \mathcal{A}} C^a$ and $D = \sum_{a \in \mathcal{A}} D^a$. We also define the maximum usable storage output O_i to be the maximum amount of stored electricity that the whole population of agents is able to discharge and make use of during time interval i . If agents can only use stored electricity to satisfy their own energy needs, then we have that $O_i = \sum_{a \in \mathcal{A}} \min(D^a, \ell_i^a)$ for all $i \in \mathcal{I}$, as each agent has no reason to discharge more energy than its current load. This is the situation that we mainly focus on in this paper. However, the analysis we present can also be applied to the case where agents can sell stored energy back to the grid, using a feed-in tariff. If we assume that agents can sell stored energy at the current grid price, then our model remains the same, except for the values of O_i which should be set equal to D , for all $i \in \mathcal{I}$. To satisfy its load profile and energy charging needs, each agent must purchase electricity from the available market, which we describe in the next subsection.

3.2 A Macro-Model of the Electricity Market

In order for widespread home energy storage to effectively help reduce peak loads and improve the overall efficiency of the electricity market, there must be some incentive for users to charge their storage devices when demand is low (or when renewable energy generators are running cheaply) and discharge them when demand is high (or when expensive peaking plants are used). As discussed in Section 2.3, both TOU and RTP pricing plans aim to incentivise users to optimise their demand by shifting to low demand periods. In particular, by sending up to date pricing information, RTP allows electricity suppliers to provide consumers with prices that accurately reflect current levels of demand and costs of generating that electricity. Hence, in our analysis, we focus on the use of RTP, and note that competition between suppliers that provide RTP pricing schemes should allow future consumers to buy electricity at or close to the current market price (Schweppe, 1988). Although the use of RTP makes the settlement process more complex than in the case of Fixed or TOU pricing (as these are independent of the reaction of the agents), and requires more interactive and intelligent energy management in the home, recent advances in smart metering technology have meant that RTP has become a plausible and attractive possibility for future smart electricity grids (as discussed in Section 2.2). Accordingly, we assume agents buy electricity under an RTP scheme where the price of electricity accurately represents the market price for power at that particular settlement period.

To this end, we consider a macro-model of the electricity market that abstracts the actual market mechanism and trading that determines the real-time price of electricity. Such an approach to modelling the supply side and the transmission system is common in the power system economics literature as it does not significantly affect the general trends observed when analysing the demand side (Kirschen & Strbac, 2004; Schweppe, 1988). Our model is based on the observation that as the aggregate demand for electricity increases, the unit price of that electricity will also increase, since more costly means of generation must be used to satisfy this additional demand. For example, Figure 2 shows the half-hour UK real-time wholesale buy prices (in the balancing market) during August and September 2009 plotted against aggregate demand.⁸ While there are numerous anomalies due to power station outages that cause short term price increases, it is clear that over a large range of demand, there is an increasing trend in the relationship between the unit cost of electricity and the total aggregate demand, with prices rising rapidly for very high demand. An analysis of the effect of these price fluctuations on a system with micro-storage is left for future work.

In several countries around the world, including the UK and the US, electricity is traded in wholesale for residential, commercial and industrial purposes in forward markets (up to several months in advance) and in the balancing market (in real-time). Within this setup, the role of suppliers is to buy electricity from the generators on the wholesale market (e.g., the National Grid runs the wholesale market in the UK and PJM runs the one on the east coast of the US) and sell retail to its consumers. In the UK, the retail market is dominated by six large suppliers that retail to 26 million residential consumers (representing around 30% of all of UK energy demand). Now, in this work, we focus exclusively on these domestic consumers since they constitute of a significant portion of the overall energy consumption

8. The data is available at www.bmreports.com.

in the UK. Yet, our approach could readily be extended to the other commercial and industrial consumers. This is in line with our main aim to show that the deployment of micro-storage within a large heterogeneous population can have significant benefits. Thus, demonstrating that we can flatten the demand of the heterogeneous domestic population implies that a similar behaviour should be achievable across the remaining 70% of the population (with even larger benefits). To this end, we model the domestic consumer market with a pricing function that reflects the expected cost that a supplier (retailing electricity to the consumers) would pay from the wholesale market. There are several advantages for a supplier to having an aggregate demand with fewer peaks from its domestic consumers. Specifically, the supplier’s exposure to peak (and typically volatile) wholesale prices is reduced and it can benefit more from long-term baseload (i.e., long-term flat load) forward contracts (Voice, Vytelingum, Ramchurn, Rogers, & Jennings, 2011). Furthermore, the peak load on transmission and distribution networks may be curtailed, reducing the need for infrastructure reinforcements, and most significantly, expensive peaking plant capacity (which may only be used for a few hours each day) can also be downsized.

Against this background, in our model, the real-time price of electricity per kWh (£/kWh) is specified by a price function $s(q) = 0.04 + 0.20(q/0.6N)^\gamma$ where q kWh is the total energy required by the $N = |\mathcal{A}|$ consumers during a half-hour time slot and γ is set to 4 to model the typical trend (i.e., monotonically and rapidly increasing) of the wholesale demand prices at which the suppliers purchase their electricity. Note that as long as the modelled demand curve is monotonically increasing to reflect the increasing unit cost of electricity with demand (which incentivises demand response), the actual demand curve is not critical to our work.

Given this function, the retail price depends on the total amount of electricity consumed by the agents, ℓ_i , and the net discharging and charging of the agents’ storage devices, b_i , such that the total amount of electricity bought from suppliers at time interval i is $q_i = \ell_i + b_i$, and the market price is given by $p_i = s(q_i)$. Hence, each agent a pays $p_i \times (\ell_i^a + b_i^a)$ and the total cost for all agents is $p_i q_i$.

3.3 Grid Performance Metrics

A key aim of this paper is to study the effect of storage on the overall system and whether the global performance of the system improves as agents adopt it. In more detail, we measure performance by considering the following standard metrics of an electricity market (Harris, 2005):

- The load factor (LF) is the average power divided by peak power, over a period of time:

$$\text{LF} = \frac{\sum_{i \in I} q_i}{|I| \times \max_{i \in I} q_i}$$

where $I \subseteq \mathcal{I}$ is a selected period of time (e.g., a day). Ideally, the LF should be at its maximum of 1 which means the aggregate load profile is completely flat. When $\text{LF} < 1$ it suggests variations in demand through each daily period. Hence, if storage strategies are effective, the LF should converge to 1 when agents are able to utilise storage to shift consumption at peak time to periods of low demand. The LF on its own only measures the aggregate load profile and does not indicate how individual

agents contribute to peaks in the system (e.g., if all agents are consuming at the same time or if only a few agents cause large peaks). To better capture such behaviours we rely on the following measure, the diversity factor.

- The diversity factor (DF) is the ratio of the sum of the individual maximum demands of various consumers of the system to the maximum demand of the complete system:

$$DF = \frac{\sum_{a \in \mathcal{A}} \max_{i \in \mathcal{I}} q_i^a}{\max_{i \in \mathcal{I}} q_i}$$

The DF is always greater than or equal to 1 and the higher its value, the less correlated the peak demands of consumers are. Less correlated peaks result in a flat aggregate profile (as peaks are interposed), and a high DF implies that users are well diversified and so will not cause large aggregate peaks (by all consuming at the same time) in the system. On the other hand, a low DF indicates that agents have more correlated load profiles which could result in peaks in demand. Now, while the LF describes the aggregate demand, the DF describes how the individual demands compare with each other and, thus, the DF can give insights into how the LF is achieved (i.e., how the individual profiles are contributing to the aggregate profiles). For example, for the same high LF, a high DF suggests well diversified profiles with individual consumer peaks spread across the day while a low DF suggests flatter individual profiles. The DF is generally useful when designing decentralised control mechanisms because it identifies the individual (rather than the aggregate) behaviours of the agents and, thus, potentially provides insights on how the individual behaviour of an agent should be modified.

- The grid carbon intensity is the amount of carbon dioxide emitted in order to deliver one unit of electricity to the consumer. It is expressed as grams of CO₂ per kWh and, ideally should be as low as possible (in order to minimise greenhouse gas emission). The carbon intensity of the UK grid for half-hourly periods for August and September 2009 is shown in Figure 3. As with the market price of the electricity, there is again a clear correlation between the carbon intensity of the electricity from the grid, and the total demand on the grid. In the UK, this is due to the use of coal-fired power stations to satisfy increasing demand, and thus, as well as reducing the total cost of electricity to the consumer, we can expect the use of micro-storage to reduce total carbon emissions by reducing the overall carbon intensity of the grid.

We note here that our pricing model depends on the aggregate demand in the domestic sector. Thus, when we report load factor and diversity factor results, we only do so for this domestic sector. However, in reality, the domestic sector is only one contributor to the total electricity demand; in the UK, it represents, as previously mentioned, approximately 30% of the total demand, with the remaining 70% being commercial and industrial consumption which remains unchanged within our model. To demonstrate the potential of domestic micro-storage to reduce UK carbon emissions from electricity generation, when we present our results regarding carbon emissions reductions, we describe the reduction of carbon emissions (as a result of micro-storage in the domestic sector) as a percentage of the




Figure 3: Half-hour UK grid carbon intensity plotted against total demand for August and September 2009 illustrating the correlation between carbon intensity and demand.

current carbon emission from all sources (i.e., domestic, commercial and industrial), assuming micro-storage in the domestic sector and no change in the remaining 70% of commercial and industrial daily profiles.

4. A Game-Theoretic Analysis of Micro-Storage

Having set the stage for the application of micro-storage in the smart grid, we now explore the theoretical underpinnings of such a system. To this end, we apply a game-theoretic framework to the models given above and characterise the resulting equilibria. These equilibria are important because they represent stable states of the system under which each agent is unable to increase its profitability by unilaterally changing its strategy. In particular, we can expect selfish agents to maintain battery usage strategies that are maximally profitable for them, and thus, these are the states that we predict will arise naturally if the market prices stabilise. We address the question of how to endow individual agents with the intelligence to reach a stable equilibrium while maximising their owner's savings later in the paper (see Section 5). For now, to ensure tractability we assume that agents have homogeneous efficiency and running costs, that is $\eta^a = \eta$ and $\mu^a = \mu$ for all $a \in \mathcal{A}$ for some η and μ . While such assumptions allow us to obtain closed-form solutions to computing the equilibria in the system, they abstract from real-world systems where agents are typically heterogeneous. However, as one of the key achievements of this work, in Section 6 we show that such assumptions do not result in a significant loss of generality and that, in fact,

our theoretical results closely approximate the empirical results we obtain in more complex environments with heterogeneous agents.

Formally, the game we consider has players which coincide with our agents, $a \in \mathcal{A}$, and the game describes the outcome of a single 24 hour interval. The pay-off an agent receives is equal to minus the total costs that it experiences when purchasing electricity that day, $-\sum_{i \in \mathcal{I}} (p_i(\ell_i^a + b_i^a) + \mu c_i^a)$. The strategy space available to each agent is the set of feasible storage profiles, $b_i^a \forall i \in \mathcal{I}$. We approximate the space of feasible aggregate storage profiles to be the set of aggregate profiles which lie between the charging and discharging limits such that the total amount of energy charged is less than or equal to the total storage capacity available to the agents and exactly equal to the total amount of energy discharged divided by the efficiency. More formally we consider the set of aggregate storage profiles $b_i \forall i \in \mathcal{I}$ such that $-O_i \leq b_i \leq C$, $\sum_{i \in \mathcal{I}} d_i = \sum_{i \in \mathcal{I}} \eta c_i$ and $\sum_{i \in \mathcal{I}} c_i \leq \kappa$, where κ is the storage capacity (see Table 1 for notation).

Requiring that the total amount charged is less than the total storage capacity is a stricter constraint than simply requiring that the capacity is never exceeded at any time. However, it is a reasonable model of storage capacity limitations for a day-long time period (given the daily-cyclic nature of demand), where demand typically goes through a single cycle of low to high to low, implying that storage devices would go through a corresponding cycle of charging to discharging to charging. Indeed, in practice, we find that at the equilibria of our simulations, there is indeed a single charging and discharging cycle during which prices cycle from low to high. Furthermore, the equilibria reached in our experiments closely agree with the equilibria predicted in this section (see Subsection 6.3).

We are also making a further approximation in considering all aggregate profiles which satisfy the aggregate capacity, charging and discharging constraints. Even if the aggregate constraints are satisfied, this does not necessarily mean that there exist feasible strategies for individual agents that give this aggregate profile. To give an example, this would be a poor model if all agents had the same values for C^a, D^a and ℓ_i^a but there was a single agent who was in possession of the majority of the storage capacity. In considering this set of aggregate profiles we are therefore assuming that storage capacity is distributed evenly amongst agents, roughly in proportion to their loads and charging and discharging capacities. This is not unreasonable given the context of the situation we are modelling. We now proceed to characterise the competitive equilibria for this game.

4.1 Competitive Equilibria as Global Optimisers

The set of competitive equilibria for the system corresponds to the set of Nash equilibria under the assumption that each individual has negligible market power. That is, we assume that each agent's electricity consumption has a negligible effect on the price of electricity, and we seek situations where no agent has an incentive to change its storage profile to reduce its cost. We choose to analyse these equilibria as they capture the steady state of a real system consisting of domestic users each owning micro-storage (optimised by an agent) and where each user's consumption has minimal effect on electricity prices (in the UK, each user represents just one home out of 26 million).

Now, suppose agents have chosen some strategy profiles, and let us consider the effect of a feasible change in strategy for one agent. That is, some agent $a \in \mathcal{A}$ considers a change

for each $i \in \mathcal{I}$, changing c_i^a to $c_i^a + \Delta c_i^a$ and d_i^a to $d_i^a + \Delta d_i^a$, giving a net change from b_i^a to $b_i^a + \Delta b_i^a$ where $\Delta b_i^a = \Delta c_i^a - \Delta d_i^a$. The change in payoff for agent a would be:

$$\sum_{i \in \mathcal{I}} (s(q_i) - s(q_i + \Delta b_i^a))(\ell_i^a + b_i^a) + (s(q_i + \Delta b_i^a)\Delta b_i^a + \mu \Delta c_i^a). \quad (1)$$

As noted in the previous section, since we are examining widespread micro-storage, we can assume that for all $i \in \mathcal{I}$ and $a \in \mathcal{A}$, ℓ_i^a , C^a , and D^a are small in comparison to ℓ_i and q_i . This is equivalent to assuming each agent has negligible market power. Thus, $s(q_i + \Delta b_i^a)$ will be very close to $s(q_i)$ and the first term in the above will be small. So, the change in payoff for agent a would be approximately:

$$\sum_{i \in \mathcal{I}} (s(q_i)\Delta b_i^a + \mu \Delta c_i^a). \quad (2)$$

This is equal to the dot product of the gradient times the vector of changes $(\Delta c, \Delta d)$, for the following function, $f(\cdot)$,⁹ which we define as,

$$f(\{c_i^a, d_i^a\}_{a \in \mathcal{A}, i \in \mathcal{I}}) = \sum_{i \in \mathcal{I}} \left(\int_0^{q_i} s(x) dx + \mu c_i \right). \quad (3)$$

Thus, the condition that each agent has no incentive to change their strategy is approximately equivalent to saying that for all $a \in \mathcal{A}$, the directional derivative of $f(\cdot)$ is positive in any direction of change that leads to a feasible storage profile. This is equivalent to saying that a vector of storage profiles is a local minimum for $f(\cdot)$ over the set of all feasible storage profiles. Since $s(\cdot)$ is increasing, $f(\cdot)$ must be convex, and since the feasible domain is closed and convex, all local minimums are also global minimums for that domain. Thus, the deterministic competitive equilibria for this game correspond to vectors of strategy profiles which minimise $f(\cdot)$ over the set of feasible strategy profiles.

These approximations are typically well suited (i.e., do not result in a significant loss of accuracy) to the large systems we consider (with millions of agents) and they have the effect of greatly reducing our search for the competitive equilibrium of a complex multi-player game to a relatively straightforward constrained optimisation problem — that of minimising $f(\cdot)$. We now proceed to find solutions to this optimisation problem.

4.2 Characterisation of the Competitive Equilibria

We can now characterise the aggregate storage profiles which form optimal solutions to the constrained optimisation problem given above. This characterisation is given in our main result, Theorem 1. However, this theorem may seem somewhat unintuitive at first. We can think of the equilibrium as being characterised by two prices, p^+ and p^- (defined below). The equilibrium strategy is then to always charge as much as possible when energy is cheap, right up to the point where the price reaches p^+ , and to always discharge when energy prices are high, right up until they fall to p^- . Thus, for any interval i , if at equilibrium, $p_i < p^+$, then all agents must be at their maximum charge rates and if $p_i > p^-$, then all agents must

9. There is no real world counterpart to $f(\cdot)$, and so little intuition for its definition – it is not necessarily equal to the total revenue, for example. It is simply a tool to help characterise global equilibrium.

be at their maximum output rates. If, at equilibrium, p_i lies strictly between p^+ and p^- , then the energy is too expensive to be worth charging, but too cheap to make discharging profitable, and so no storage activity occurs for that time interval.

In order to formally state and prove this result, we must begin with a definition:

Definition 1. For a storage system as described, for each interval i , let us define the functions $\Phi_i(\cdot)$ and $\Psi_i(\cdot)$ to be,

$$\begin{aligned}\Phi_i(p) &= \min(C, (s^{-1}(p) - \ell_i)^+), \\ \Psi_i(p) &= \min(O_i, (s^{-1}(p) - \ell_i)^-).\end{aligned}$$

That is, $\Phi_i(p)$ is the amount of electricity that would have to be charged during interval i , if there were no discharging, in order for the resulting price to be as close to p as possible. Similarly, $\Psi_i(p)$ is the amount of electricity that would have to be discharged during interval i , if there were no charging, in order for the resulting price to be as close to p as possible. O_i is the maximum useful discharge during i and C , the maximum daily total charge.

We define the *discharging price point*, p^- , to be the maximum of the union of the solutions to:

$$\sum_{i \in \mathcal{I}} \Psi_i(p^-) = \eta \sum_{i \in \mathcal{I}} \Phi_i(\eta p^- - \mu),$$

and the solutions to:

$$\sum_{i \in \mathcal{I}} \Psi_i(p^-) = \eta \kappa,$$

if such exist. This maximum exists because $s(\cdot)$ is continuous and strictly increasing.

If p^- is well defined, then we also define the *charging price point*, p^+ , to be $\eta p^- - \mu$ if

$$\sum_{i \in \mathcal{I}} \Psi_i(p^-) < \eta \kappa,$$

or equal to the minimal solution to:

$$\sum_{i \in \mathcal{I}} \Phi_i(p^+) = \kappa,$$

if such exist, otherwise.

Note, if p^- and p^+ are well defined then either $p^+ = \eta p^- - \mu$ and $\sum_{i \in \mathcal{I}} \Psi_i(p^-) = \eta \sum_{i \in \mathcal{I}} \Phi_i(p^+)$, or $\sum_{i \in \mathcal{I}} \Psi_i(p^-) = \eta \kappa$, and $\sum_{i \in \mathcal{I}} \Phi_i(p^+) = \kappa$. In the latter case, since p^+ is strictly greater than all solutions to $\sum_{i \in \mathcal{I}} \Psi_i(p^-) = \eta \sum_{i \in \mathcal{I}} \Phi_i(\eta p^- - \mu)$, and since the $\Phi_i(\cdot)$ functions are increasing and the $\Psi_i(\cdot)$ functions are decreasing, we must have that:

$$\eta \sum_{i \in \mathcal{I}} \Phi_i(p^+) = \eta \kappa = \sum_{i \in \mathcal{I}} \Psi_i(p^-) < \eta \sum_{i \in \mathcal{I}} \Phi_i(\eta p^- - \mu),$$

and so $p^+ < \eta p^- - \mu$. Under our intuitive understanding of the equilibrium given above, this means that, p^- and p^+ are chosen so that the total amount discharged is equal to the total amount charged times η . Furthermore, under this restriction, either p^- and p^+ are chosen to maximise the total amount charged or else the total amount charged is equal to κ .

Lemma 1. *There always exists a solution to $\sum_{i \in \mathcal{I}} \Psi_i(p) = \eta \sum_{i \in \mathcal{I}} \Phi_i(\eta p - \mu)$. Furthermore, if p is the maximal solution then $p^- = p$ and $p^+ = \eta p - \mu$ unless $\sum_{i \in \mathcal{I}} \Psi_i(p) > \eta \kappa$, in which case, p^- is the maximal solution to $\sum_{i \in \mathcal{I}} \Psi_i(p^-) = \eta \kappa$, and p^+ is the minimal solution to $\sum_{i \in \mathcal{I}} \Phi_i(p^+) = \kappa$.*

Proof. We have that $s(\cdot)$ is a continuous, strictly increasing function, and $\ell_i > 0$ is inside its range for each time interval i . Thus, if p is sufficiently small, then for all $i \in \mathcal{I}$ we will have $s^{-1}(p) < \ell_i$ and hence $\Psi_i(p)$ will be strictly positive and, since $\eta p - \mu < p$, $\Phi_i(\eta p - \mu)$ will be zero. Likewise if p is sufficiently large then for all $i \in \mathcal{I}$ we will have $s^{-1}(p) > \ell_i$ and so $\Phi_i(p)$ will be strictly positive and $\Psi_i((p + \mu)/\eta)$ will be zero. Since the functions $\Psi_i(\cdot)$ are decreasing for all i and $\Phi_i(\cdot)$ are increasing for all i , we can conclude that $\sum_{i \in \mathcal{I}} \Psi_i(p) - \eta \sum_{i \in \mathcal{I}} \Phi_i(\eta p - \mu)$ is a continuous decreasing function in p which is negative for sufficiently small p and positive for sufficiently large p . This implies the existence of some solution p such that $\sum_{i \in \mathcal{I}} \Psi_i(p) = \eta \sum_{i \in \mathcal{I}} \Phi_i(\eta p - \mu)$. Let p be the maximal solution to this.

Now if $\sum_{i \in \mathcal{I}} \Psi_i(p) \leq \eta \kappa$ then, since the $\Psi_i(\cdot)$ functions are decreasing, there can be no $p' > p$ such that $\sum_{i \in \mathcal{I}} \Psi_i(p') = \eta \kappa$. Thus, $p^- = p$ and, so, $p^+ = \eta p - \mu = \eta p^- - \mu$. If $\sum_{i \in \mathcal{I}} \Psi_i(p) > \eta \kappa$ then, since $\sum_{i \in \mathcal{I}} \Psi_i(\cdot)$ is a decreasing function that eventually reaches zero, it must cross $\eta \kappa$. Thus, there will be a maximal value for p^- such that $\sum_{i \in \mathcal{I}} \Psi_i(p^-) = \eta \kappa$. Similarly, since $\sum_{i \in \mathcal{I}} \Phi_i(p) > \kappa$, there must exist a minimal value for p^+ such that $\sum_{i \in \mathcal{I}} \Phi_i(p^+) = \kappa$, as required. \square

Although we have specified p^- so that it is the maximum solution to $\sum_{i \in \mathcal{I}} \Psi_i(p^-) = \eta \sum_{i \in \mathcal{I}} \Phi_i(\eta p^- - \mu)$, if $\sum_{i \in \mathcal{I}} \Psi_i(p^-) < \eta \kappa$, it is worth noting that for any two values p and p' that satisfy $\sum_{i \in \mathcal{I}} \Psi_i(p) = \eta \sum_{i \in \mathcal{I}} \Phi_i(\eta p - \mu)$ and $\sum_{i \in \mathcal{I}} \Psi_i(p') = \eta \sum_{i \in \mathcal{I}} \Phi_i(\eta p' - \mu)$, by the monotonicity on both sides of the equation, we must have that for all $i \in \mathcal{I}$, $\Psi_i(p) = \Psi_i(p')$ and $\Phi_i(p) = \Phi_i(p')$. Likewise, if $\sum_{i \in \mathcal{I}} \Psi_i(p) = \sum_{i \in \mathcal{I}} \Psi_i(p')$ or $\sum_{i \in \mathcal{I}} \Phi_i(p) = \sum_{i \in \mathcal{I}} \Phi_i(p')$ then for all $i \in \mathcal{I}$, $\Psi_i(p) = \Psi_i(p')$ or $\Phi_i(p) = \Phi_i(p')$ respectively. Indeed, with these observations in mind, the specifications that maximal solutions be chosen makes no difference to our main result, and is done simply so that p^+ and p^- are well defined.

We can now state and prove the main result of this analysis.

Theorem 1. *For a storage system as described, if $\eta < 1$ or $\mu > 0$ then the set of competitive equilibria for the system is precisely the set of feasible agent strategies where, for all $i \in \mathcal{I}$, $c_i = \Phi_i(p^+)$ and $d_i = \Psi_i(p^-)$. In the case where storage devices are perfectly efficient and costless, ($\eta = 1$ and $\mu = 0$), the set of competitive equilibria for the system is precisely the set of feasible agent strategies where, for all $i \in \mathcal{I}$, $b_i = \Phi_i(p^+) - \Psi_i(p^-)$.*

Proof. We seek to find an aggregate storage profile, which minimises $f(\cdot)$ where:

$$f(\{c_i, d_i\}_{i \in \mathcal{I}}) = \sum_{i \in \mathcal{I}} \left(\int_0^{\ell_i + c_i - d_i} s(x) dx + \mu c_i \right).$$

Since the set of feasible aggregate storage profiles is closed and bounded, there must be at least one minimum of $f(\cdot)$ over this domain. However, since this domain is convex, and $f(\cdot)$ is a convex function, the only local minima will be a convex set of global minima. To find these optimal allocations, we seek feasible aggregate storage profiles for which the

derivative of $f(\cdot)$ is non-negative in every direction that leads to another feasible allocation. We can calculate that for all $i \in \mathcal{I}$, $\partial f/\partial c_i = p_i + \mu$ and $\partial f/\partial d_i = p_i$. Thus it remains to characterise all feasible profiles such that:

$$\sum_{i \in \mathcal{I}} (p_i + \mu)\Delta c_i + p_i\Delta d_i \geq 0,$$

for every vector of small changes that lead to another feasible aggregate storage profile, that is, all $\Delta c, \Delta d$ such that $\{c_i + \Delta c_i, d_i + \Delta d_i\}_{i \in \mathcal{I}}$ is feasible.

Now suppose we have some storage profile, (c, d) , which locally minimises $f(\cdot)$. If there are time intervals i, j with $C > c_i > 0$ and $C > c_j > 0$, then it would be feasible to increase c_i and decrease c_j by an equal quantity, (or vice versa), hence we must have $p_i = p_j$. From this we can deduce that if for some i, j , $p_i < p_j$, $c_i > 0$ and $c_j > 0$, then we must have $c_i = C$. So, if we let \hat{p}^+ be the maximum of p_j for time intervals j with $c_j > 0$, we get that for all intervals i such that $c_i > 0$, $p_i \leq \hat{p}^+$, and if $p_i < \hat{p}^+$, then c_i must equal C . Similarly, we can show that if \hat{p}^- is the minimum of p_j for time intervals j such that $d_j > 0$ then for all intervals i that $d_i > 0$, $p_i \geq \hat{p}^-$, and if $p_i > \hat{p}^-$ then $d_i = 0$.

Furthermore, there cannot be i, j such that $p_i + \mu > \eta p_j$ and $c_i > 0$ and $d_j > 0$, for then it would be feasible and profitable to decrease c_i by some Δc_i and increase d_j by $\Delta d_j = \eta \Delta c_i$. Hence, we must have $\hat{p}^+ + \mu \leq \eta \hat{p}^-$. In particular, $\hat{p}^+ \leq \hat{p}^-$, with this inequality being strict if $\mu > 0$ or $\eta < 1$, in which case, for all intervals $i \in \mathcal{I}$ at most one of c_i and d_i can be non-zero.

For each interval i , if $p_i > \hat{p}^-$, then $d_i = 0$ and $c_i = 0$, so $s^{-1}(p_i) - \ell_i = -O_i$ and thus, $s^{-1}(\hat{p}^-) - \ell_i < -O_i$ as $s^{-1}(\cdot)$ is an increasing function. This means that $b_i = -O_i = -\Psi_i(\hat{p}^-)$. On the other hand, if $p_i = \hat{p}^-$ then $b_i = s^{-1}(\hat{p}^-) - \ell_i = -\Psi_i(\hat{p}^-)$. Likewise, for each interval i , if $p_i < \hat{p}^+$, then $c_i = C$ and $d_i = 0$, and so $s^{-1}(p_i) - \ell_i = C$ and so $\Phi(\hat{p}^+) = C = b_i$, and if $p_i = \hat{p}^+$ then $b_i = \Phi_i(\hat{p}^+)$, directly. If p_i lies strictly between \hat{p}^- and \hat{p}^+ then both c_i and d_i must equal 0, and so $s^{-1}(p_i) - \ell_i = 0$. Thus, $s^{-1}(\hat{p}^-) - \ell_i \geq 0$ and $s^{-1}(\hat{p}^+) - \ell_i \leq 0$, and so $\Phi_i(\hat{p}^+) = \Psi_i(\hat{p}^-) = 0$. These results show that for all prices, at equilibrium we must have $b_i = \Phi_i(\hat{p}^+) - \Psi_i(\hat{p}^-)$. As stated above, if $\eta < 1$ or $\mu > 0$, then for each time interval i , at most one of c_i and d_i can be non-zero, and so, in this case we can specify, $c_i = \Phi_i(\hat{p}^+)$ and $d_i = \Psi_i(\hat{p}^-)$.

Since $\sum_{i \in \mathcal{I}} \Psi_i(\hat{p}^-) = \eta \sum_{i \in \mathcal{I}} \Phi_i(\hat{p}^+)$ and $\hat{p}^+ + \mu \leq \eta \hat{p}^-$, we must have that $\sum_{i \in \mathcal{I}} \Psi_i(\hat{p}^-) \leq \eta \sum_{i \in \mathcal{I}} \Phi_i(\eta \hat{p}^- - \mu)$ and so, any solution to $\sum_{i \in \mathcal{I}} \Psi_i(p) = \eta \sum_{i \in \mathcal{I}} \Phi_i(\eta p - \mu)$ must satisfy $p \leq \hat{p}^-$, by the monotonicity of both sides of this equation. Furthermore, for any such p , we would have $\sum_{i \in \mathcal{I}} \Psi_i(p) \geq \sum_{i \in \mathcal{I}} \Psi_i(\hat{p}^-)$, meaning that $\sum_{i \in \mathcal{I}} \Phi_i(\eta p - \mu) \geq \sum_{i \in \mathcal{I}} \Phi_i(\hat{p}^+)$ which implies that $\eta p - \mu \leq \hat{p}^+$. It would also mean that, since $\sum_{i \in \mathcal{I}} d_i = \sum_{i \in \mathcal{I}} \Psi_i(\hat{p}^-) \leq \eta \kappa$, by the capacity constraint, we must have $\sum_{i \in \mathcal{I}} \Psi_i(\hat{p}^-) \leq \sum_{i \in \mathcal{I}} \Psi_i(p^-)$. Similarly, we must also have that $\sum_{i \in \mathcal{I}} \Phi_i(\hat{p}^+) \leq \sum_{i \in \mathcal{I}} \Phi_i(p^+)$.

Now, suppose that $\sum_{i \in \mathcal{I}} \Phi_i(\hat{p}^+) < \sum_{i \in \mathcal{I}} \Phi_i(p^+)$. This would also imply that $\sum_{i \in \mathcal{I}} \Psi_i(\hat{p}^-) < \sum_{i \in \mathcal{I}} \Psi_i(p^-)$. Hence, there would have to exist some intervals i, j with $s^{-1}(\hat{p}^+) \leq \ell_i < s^{-1}(p^+)$ and $s^{-1}(\hat{p}^-) \geq \ell_j > s^{-1}(p^-)$. However, that would imply that, in our equilibrium state, $c_i = d_i = 0$, $d_j = c_j = 0$. Hence we would have $p_i < p^+$ and $p_j > p^-$, meaning that $p_i + \mu < \eta p_j$ and that it would be profitable to increase c_i by some Δc_i and increase d_j by $\Delta d_j = \eta \Delta c_i$. This implies a contradiction, since for a small enough change, this would lead to a feasible storage profile.

Hence we must have that $\sum_{i \in \mathcal{I}} \Phi_i(\hat{p}^+) = \sum_{i \in \mathcal{I}} \Phi_i(p^+)$ and $\sum_{i \in \mathcal{I}} \Psi_i(\hat{p}^-) = \sum_{i \in \mathcal{I}} \Psi_i(p^-)$. However, since these functions are monotonic, we must have that for all i , $\Phi_i(\hat{p}^+) = \Phi_i(p^+)$ and $\Psi_i(\hat{p}^-) = \Psi_i(p^-)$. Thus, our equilibrium must satisfy the conditions in the statement of this theorem.

Thus, there exists at least one minimiser of $f(\cdot)$ over the feasible domain, and any minimiser of $f(\cdot)$ must satisfy the conditions given in the statement of the theorem. Furthermore, the conditions given in the statement of the theorem are sufficiently strict as to specify $f(\cdot)$ precisely. Thus, any feasible storage profile which satisfies the conditions given in the statement of the theorem must minimise $f(\cdot)$ over the feasible domain, as required. \square

A direct consequence of this theorem and the prior observations is that no matter what the storage capacity is, the aggregate amount charged during the day is bounded above by:

$$\sum_{i \in \mathcal{I}} \Phi_i(\eta p - \mu),$$

for p equal to any solution to:

$$\sum_{i \in \mathcal{I}} \Psi_i(p) = \eta \sum_{i \in \mathcal{I}} \Phi_i(\eta p - \mu).$$

If the storage capacity is greater than this amount, then some portion of it will not be used at equilibrium. Moreover, since we characterise the competitive equilibria as global minimisers of aggregate costs, and agents have negligible market power, the addition of more storage capacity is profitable if and only if the total amount of storage is less than this maximal value. This key result leads us to predict that for a given aggregate load profile, either not all consumers will need to buy storage or the optimal battery capacity to buy for each consumer will be bounded. In particular, this bound is further supported by our empirical evaluation of the UK electricity market where only a subsection of the population is required to adopt storage to minimise their costs (see Section 6).

4.3 Idealised Scenarios

Having determined the existence and characterisation of charging and discharging price points, we now investigate how these prices will be set in the context of two idealised scenarios where micro-storage devices have been deployed in the grid on a large scale. This aims to identify the properties of the system as different parameters tend to particular limits (and understand the broad system behaviour). Our intuitions on their impact are expressed through the following corollaries. We consider the situation where agents can sell electricity back to the grid in order to simplify the results we obtain and thus make clearer the intuition we are trying to provide. Similar results would hold in the case where agents cannot sell, but have similar-shaped load profiles, and storage capacities in proportion to their daily load.

Corollary 1. *If agents are allowed to sell electricity back to the grid at the current price, and if, for all agents $a \in \mathcal{A}$, C^a and D^a are sufficiently large, then for all i , $p^+ \leq p_i \leq p^-$. Furthermore, if for any $i \in \mathcal{I}$, $p^+ < p_i < p^-$, then $b_i = 0$.*

Proof. If, for all $a \in \mathcal{A}$ we let C^a and D^a be equal to κ^a , then this does not break our smallness assumption, and, furthermore, for all $i \in \mathcal{I}$, we'll have $\Psi_i(p^-) < O_i$ and $\Phi_i(p^+) < C$. For all i , either $b_i < 0$, in which case $0 < \Psi_i(p^-) < D$ and so $p_i = p^-$, or else $b_i > 0$, in which case $0 < \Phi_i(p^+) < C$ and so $p_i = p^+$, or, lastly, we could have $b_i = 0$, which can only occur if $\Phi_i(p^+) = \Psi_i(p^-)$, and so either $p^+ = p_i = p^-$ or $\Phi_i(p^+) = \Psi_i(p^-) = 0$ and $p^+ < p_i < p^-$. This covers all possible cases, as required. \square

Thus, if the charge and discharge rates are sufficiently high, then we could expect prices to always lie within p^+ and p^- .

Corollary 2. *If agents are allowed to sell electricity back to the grid at the current price, and if, for all agents a , C^a and D^a are sufficiently large, capacity κ is sufficiently high, $\mu = 0$ and $\eta = 1$, then for all i , $p^+ = p_i = p^-$.*

Proof. If κ is sufficiently high, then we must have an equilibrium where $p^+ = \eta p^- - \mu$, or $p^+ = p^-$. The result then follows directly from the previous corollary. \square

Hence, in the above scenario, with perfectly efficient, cost free, and high capacity storage, we would expect the market prices over time to flatten to a single value. This is because perfect storage capability would allow agents to transport energy from any time interval to any other time interval free of charge. Thus, different suppliers in different time intervals would have to compete with each other, resulting in convergence to a single market price. Even if storage devices are not perfectly efficient, they still have a price flattening effect, though this is mitigated by the fact that agents would have to buy more energy than they discharge, meaning they would always require some price difference between charging and discharging periods. Since prices are a direct function of demand, we can infer that having large amounts of storage should effectively bound the maximum and minimum levels of demand for each time period. Thus, we should expect the addition of large amounts of storage to have a significant effect on the grid load factor, with the size of this effect directly related to the efficiency of the storage device. Hence, in Section 6, we study the impact of storage adoption on the grid (i.e., in terms of grid efficiency and cost savings for the population) and determine the impact of micro-storage devices for various levels of saturation of micro-storage across the population. Before doing so, however, in the next section, we provide an analysis of the proportion of the population expected to adopt storage devices (i.e., based on an individual's profits from micro-storage). This is important because, having shown that the real-time price of electricity should flatten, we aim to show how such prices will impact on the profitability of buying micro-storage. Such an analysis can then be helpful in determining the viability of micro-storage deployment projects.

4.4 Storage Adoption

Our model shows that as available storage capacity and charging (and discharging) rates increase, there can be a significant effect on prices. Moreover, it shows there is *a limit on how much storage capacity is required*, such that it is profitable to increase capacity if and only if the available storage is below this limit. Finding such limits is a useful application of this analysis, as it gives an indication of what level of adoption is likely to occur given the cost, efficiency and charging and discharging rates of the leading storage technology.

In reality, however, it is not practical for home users to incrementally increase their storage capacity over time to find the optimal level. The more likely scenario is that when a home user buys a storage device, they will buy enough to cover their usage requirements. Aggregate storage capacity will increase over time as more and more homes install such devices. Predicting the number of potential buyers of micro-storage devices will be key to understanding (for market makers and producers) whether demand will be large enough to take advantage of the economies of scale for the production of micro-storage devices (of high efficiency or high capacity). In this respect, it is crucial to study the maximal level of storage adoption for such batteries. That is, the percentage of homes that can install electricity storage devices before it is no longer profitable for more devices to be installed.

We will consider the case where home users cannot sell electricity back to the grid (or to neighbours). If users can sell stored electricity, then it would be possible that some users could purchase extra capacity or multiple storage devices as a way to make money. However, the devices themselves are likely to be manufactured with the energy needs of a single home in mind, and so this analysis still gives a useful guide to how many such devices can be profitably deployed throughout the populace.

Corollary 3. *Suppose we model a population of agents \mathcal{A} such that a subset of agents $\mathcal{A}' \subset \mathcal{A}$ have homogeneous storage devices, where $|\mathcal{A}'| = \rho|\mathcal{A}|$, for some $\rho < 1$. Suppose further that the population of agents in \mathcal{A}' is homogeneous, so that the aggregate load profile of agents in \mathcal{A}' is $\rho\ell_i$ and that storage devices have sufficiently large maximal charging and discharging rates and capacities so that they do not restrict their storage profiles at equilibrium. Then, at equilibrium, it is individually profitable and increases aggregate benefit for an agent in $\mathcal{A} \setminus \mathcal{A}'$ to install a storage device if and only if:*

$$\rho < \max_{i \in \mathcal{I}} \frac{1}{\ell_i} (s^{-1}(p) - \ell_i)^-,$$

where p is the maximal solution to:

$$\sum_{i \in \mathcal{I}} (s^{-1}(\eta p - \mu) - \ell_i)^+ = \sum_{i \in \mathcal{I}} \min(\ell_i, (s^{-1}(p) - \ell_i)^-).$$

Proof. We can model this scenario by setting:

$$D^a = C^a = \max_{i \in \mathcal{I}} \ell_i^a,$$

and

$$\kappa^a = \sum_{i \in \mathcal{I}} \ell_i^a,$$

for all $a \in \mathcal{A}'$ and $C^a = C^a = \kappa^a = 0$ for all $a \in \mathcal{A} \setminus \mathcal{A}'$. For the agents in \mathcal{A}' , these values for maximal charging and discharging rates and capacities are sufficiently large that the corresponding constraints cannot be tight at equilibrium.

From Theorem 1, the aggregate storage profile at equilibrium will be:

$$(s^{-1}(\eta p^- - \mu) - \ell_i)^+ - \min(\rho\ell_i, (s^{-1}(p^-) - \ell_i)^-)$$

where p^- is the maximal solution to:

$$\sum_{i \in \mathcal{I}} (s^{-1}(\eta p^- - \mu) - \ell_i)^+ = \sum_{i \in \mathcal{I}} \min(\rho \ell_i, (s^{-1}(p^-) - \ell_i)^-). \quad (4)$$

It is profitable for an agent in $\mathcal{A} \setminus \mathcal{A}'$ to get a storage device if and only if the addition of that storage device will increase the total amount of energy charged and discharged at equilibrium. For if the addition of the storage device does lead to an increase in the total amount of energy charged, it means that it is more profitable for the agent to use its storage device than not to use it. Thus, the agent must obtain some profit by having the device. If the addition of the storage device would not lead to an increase in the total amount of energy charged, then, since the amount charged is:

$$\sum_{i \in \mathcal{I}} \min(\rho \ell_i, (s^{-1}(p^-) - \ell_i)^-),$$

This means the addition of the storage device can have no effect on p^- . Hence, in this circumstance, the addition of a storage device would have no effect on the aggregate storage profile, and so the collection of storage profiles that agents have in equilibrium would still remain an equilibrium if a new device was added – with the maximally profitable behaviour of the agent with the new device being to simply not use it.

If for any $i \in \mathcal{I}$,

$$\rho \ell_i < (s^{-1}(p^-) - \ell_i)^-,$$

then, since the addition of a storage device will strictly increase ρ , the previous value of p^- will no longer satisfy (4) and so p^- , along with the total amount of energy charged, will increase. Otherwise, increasing ρ has no effect on (4), meaning p^- will not change, and neither will the aggregate amount of energy charged.

Let p be the maximal solution to:

$$\sum_{i \in \mathcal{I}} (s^{-1}(\eta p - \mu) - \ell_i)^+ = \sum_{i \in \mathcal{I}} \min(\ell_i, (s^{-1}(p) - \ell_i)^-).$$

If

$$\rho \ell_i < (s^{-1}(p^-) - \ell_i)^-$$

for some $i \in \mathcal{I}$ then, since, by inspection $p \geq p^-$, we must have:

$$\rho \ell_i < (s^{-1}(p) - \ell_i)^-.$$

If

$$\rho \ell_i \geq (s^{-1}(p^-) - \ell_i)^-$$

for all i , then p^- satisfies:

$$\sum_{i \in \mathcal{I}} (s^{-1}(\eta p^- - \mu) - \ell_i)^+ = \sum_{i \in \mathcal{I}} (s^{-1}(p^-) - \ell_i)^-,$$

and, since $\ell_i > \rho \ell_i$, p^- satisfies:

$$\sum_{i \in \mathcal{I}} (s^{-1}(\eta p^- - \mu) - \ell_i)^+ = \sum_{i \in \mathcal{I}} \min(\ell_i, (s^{-1}(p^-) - \ell_i)^-),$$

and so, $p^- = p$ and:

$$\rho \ell_i \geq (s^{-1}(p) - \ell_i)^-.$$

Thus, the condition for an additional storage device to be profitable is that for some $i \in \mathcal{I}$,

$$\rho \ell_i < (s^{-1}(p) - \ell_i)^-$$

as required. \square

This corollary can be used to give an indication of the level of adoption in a population required to see maximal aggregate cost savings from the use of energy storage. Later, in Sections 6.6 and 7, we complement this result with an empirical study of the system-wide benefits of micro-storage adoption in the UK market (where agents use our novel storage strategies) which points to a similar bound on the level of storage adoption.

4.5 Rationality Assumption

So far, our main results (i.e., Theorem 1 and Corollaries 1, 2, and 3) give the aggregate storage behaviour when our game is in a deterministic competitive equilibrium and predict the extent and nature of the adoption of storage devices in a population. We can use these results to specify the limits of the grid performance benefits and market conditions (i.e., levels of adoption and equilibrium price for electricity) that can result from adopting micro-storage. If the actions of such selfish and profit-motivated agents are to result in stable aggregate behaviour, then we can do no better than the outcomes described above.

However, in using game theory, we have made some implicit assumptions, specifically that agents are rational and have complete information about the market throughout the time period and have the ability to compute an optimal strategy given that information. In reality, information available to those owning storage devices will not be perfect and the agents will have different computational capabilities. Furthermore, even with perfect information, it might not be apparent to an automated agent which strategies are preferable. Instead, the agents themselves must adapt over time, to become aware of any changes in market prices, and learn which storage strategies are preferable. This is a difficult problem and it is not guaranteed that selfish learning behaviour can converge. In particular, if agents over-react to perceived opportunities in the market, cycles of price fluctuations could develop with no stable outcome (as seen in Subsection 6.2).

Having looked at an analytical approach which required complete information, in the next section, we provide a practical and (informationally) decentralised approach that addresses this lack of complete information. Specifically, we describe a novel adaptive storage strategy that dynamically adapts to changes in market prices, allowing the selfish, profit-motivated agents to individually maximise their savings using only their private information and information about observed market prices. Under this scheme, agents learn to change their storage profiles each day to be closer to their perceived optimal strategy. In Section 6, we show that provided the adaptation is not too fast, our mechanism converges to an equilibrium predicted by Theorem 1. Moreover, our empirical results confirm the bounds on storage capacity and adoption of storage we have predicted so far. Altogether, these results show that the assumptions we have made in our theoretical framework are reasonable

enough to model heterogeneous populations of agents and, therefore, our framework can be generally applied to large-scale micro-storage analysis.

5. An Adaptive Storage Strategy

In this section, we present a novel adaptive strategy that an agent can use to decide when to store energy and when to use the energy it has stored. Now, for the system to converge towards an equilibrium, we need to avoid having too many agents charging their batteries at the same time, in turn resulting in higher costs for everyone. Any strategy that achieves this would be a good candidate as long as it is shown to converge to the equilibrium. One possible candidate would allow the agents to adapt their storage profile solely using a target profile (which would be the equilibrium profile in this case) provided by the supplier (similar to existing TOU pricing schemes which, rather poorly, incentivise charging during off-peak hours and discharging at peak time). However, such a strategy would require optimisation by a centre (i.e., the supplier) and the solution would depend on how accurately the supplier can estimate the combined charging and discharging rate limits and storage capacities of all agents with storage devices across the grid and how often it needs to do so. Thus, we prefer a strategy that does not require grid-wide knowledge and that can adapt based only on the agents' private information (i.e., information about their own micro-storage devices) and the observed market information (i.e., real-time retail prices that continuously change as a result of changing demand due to consumers using storage devices). Indeed, we design such a decentralised strategy which we now describe in more detail.

Our strategy is based on a day-ahead best-response storage.¹⁰ Because the market prices are unknown *a priori*, we can only calculate the storage profile on a day-ahead basis, as a best-response to the predicted market prices (which are only observed *a posteriori* once the aggregate demand of the market is known). Now, if all agents were to adopt their best-response, the resulting effect would simply be peaks moving from periods of high demand to those with previously low demand as we empirically demonstrate in Subsection 6.2. With the peaks in demand moved to previously low-demand periods (as are peaks in market prices), the agents end up paying higher prices when they charge their battery. Thus, an agent that plays its best-response is exposed to these changing peaks. To mitigate its exposure to these changes, we need to ensure that each agent *gradually* adapts its storage towards the best-response storage instead of *reacting* to prices and, by so doing, avoid all agents herding to consume at the lowest predicted price point. In this section, we first describe how we calculate the day-ahead best response storage profile and, second, we describe our learning mechanism, that is how the agent adapts its storage.

10. Our strategy is unaffected by the use of different time-scales (other than day-ahead). We perform the day-ahead optimisation in our case as there exists a natural cycle of consumption over similar days that we can exploit to generate load profile distributions in our simulations and reduce the number of times the optimisation algorithms have to be run. In a real-world deployment, a finer grained optimisation (at the level of half or quarter hours) would probably be more appropriate as the agent re-optimises based on the up-to-date intra-day half-hourly consumption and market prices as well as the current amount of stored energy available.

5.1 The Day-Ahead Best-Response Storage

The objective of agent a is to minimise its costs by storing energy when prices are low and using that energy when prices are high. Now, because market prices are unknown until the aggregated load of all consumers $\sum_{a \in \mathcal{A}} \ell_i^a = \ell_i$ is known, the agent needs to decide on its storage profiles based on a prediction of the market prices. Note that in our work, we assume that market prices do not move significantly over similar days (e.g., during the same season, weekdays tend to be similar to each other but different to week ends) and use a weighted moving average to predict these future market prices.¹¹

We compute the storage profile, $b_i^a = c_i^a - d_i^a$ at every time-slot during the day as the solution to the optimisation problem (expressed as a linear program) where we minimise the following cost function given the decision variables c_i^a, d_i^a , and $\tilde{\kappa}^a$ (representing the storage capacity):¹²

$$\arg \min_{b^a} \left(\sum_{i \in \mathcal{I}} p_i (c_i^a - d_i^a + \ell_i^a) + \mu^a c_i^a \right) \quad (5)$$

subject to the following constraints:

Constraint 1: Energy conservation

$$\sum_{i \in \mathcal{I}} d_i^a = \eta^a \sum_{i \in \mathcal{I}} c_i^a$$

Constraint 2: Within charging and discharging rate limits

$$d_i^a \leq D^a \text{ and } c_i^a \leq C^a \quad \forall i \in \mathcal{I}$$

Constraint 3: Energy that can be stored or used at a time-slot

$$\begin{aligned} d_i^a &\leq \eta^a \left(c_0^a + \sum_{j=1}^{i-1} (c_j^a - d_j^a) \right) \quad \forall i \in \mathcal{I} \\ c_i^a &\leq \tilde{\kappa}^a - \left(c_0^a + \sum_{j=1}^{i-1} (c_j^a - d_j^a) \right) \quad \forall i \in \mathcal{I} \end{aligned}$$

Constraint 4: No reselling allowed

$$\ell_i^a - d_i^a \geq 0 \quad \forall i \in \mathcal{I}$$

The last constraint can be removed in a system where consumers are allowed to sell power to the grid. Note that when the capacity of the storage κ^a is known, $\tilde{\kappa}^a$ is constrained to κ^a and when we need to find κ^a , $\tilde{\kappa}^a$ is left unconstrained in the optimisation and κ^a is then calculated as the maximum energy stored (in the optimised storage profile), i.e., $\max_{i \in \mathcal{I}} \left(c_0^a + \sum_{j=1}^i b_j^a \right)$.

11. As we will demonstrate later on, this is not central to our work as the price movements are generally small. However, a number of more sophisticated prediction algorithms, such as regression or Gaussian processes, could be used instead for better predictions if price movements were significant as a result of a very volatile market.

12. We used IBM ILOG CPLEX 12 to implement and solve the linear program.

As described earlier, μ^a is the cost of using storage which we set to be very small as we wish to find the best response regardless of the external factor which is the cost of storage. η^a is the efficiency of the storage device, C^a is the maximum charging rate, D^a the maximum discharging rate and c_0^a is the amount of energy stored at the beginning of the day which is equal to the stored energy at the end of the day. We next consider how the agent adapts its storage based on its best-response.

5.2 Learning in the Market

Because market prices move each day, the agent needs to continuously adapt its storage profile to reflect these changes. One may expect that if micro-storage is incrementally rolled out, the agents would be able to gradually adapt and stabilise market prices. However, as we will empirically demonstrate in Subsection 6.2, the system becomes unstable when too many agents attempt to optimise at the same time, even if they use their best response and incrementally acquire micro-storage devices. Now, because of the relatively high cost of storage (compared to the savings in energy cost — see Section 7) and assuming the effect of micro-storage in the system will change market prices significantly (such that the optimal capacity the agent requires will be changing — as predicted by the results in Section 4.2), it is necessary for the agent to gradually change how much of its absolute storage capacity that it actually uses. To this end, based on intuitions drawn from our analytical results that point to a bound on the capacity required and the adaptation of charging profiles to prices at different times of the day, we develop a learning mechanism based on the Widrow-Hoff learning rule (i.e., a gradient descent approach)¹³ that adapts both the storage profile and capacity of the micro-storage device that is used with respect to changes in the market prices.

Our learning mechanism is based on a two-pass approach to adapt the storage capacity and profile. Initially, the agent computes the optimal storage capacity \varkappa^a required to minimise its costs. More precisely, \varkappa^a is the cost-minimising capacity by setting $\tilde{\kappa}^a$ as an unconstrained variable in the optimisation (see Equation (5)). Now, \varkappa^a constitutes a *desired capacity* towards which the agent adapts its utilised storage capacity. That is, it changes its storage capacity that it uses progressively to follow the changing trend of market prices. The actual storage capacity used by the agent is defined by Equation (6) as $\kappa^a(t)$ that follows the desired storage capacity \varkappa^a such that:

$$\kappa^a(t+1) = \kappa^a(t) + \gamma_1(\varkappa^a - \kappa^a(t)) \quad (6)$$

where $\kappa^a(0) = 0$ by default and γ_1 is the learning rate of the storage capacity of agent a .¹⁴ Given its storage capacity, the agent then computes its optimal storage profile for the following day by *fixing* κ^a at $\kappa^a(t+1)$ in Equation (5).

On the second pass, given its current storage profile, the agent adapts its storage profile as follows:

13. We used the Widrow-Hoff rule as it can be directly engineered into our optimisation algorithm, but with relatively minor changes, other learning rules (such as reinforcement learning or Bayesian learning) could be used.

14. As we will empirically demonstrate in Section 6, the choice of the learning rates determines the evolutionary stability of the system and has to be reasonably small to ensure convergence.

$$b_i^a(t+1) = b_i^a(t) + \gamma_2(\beta_i^a - b_i^a(t)) \quad \forall i \in \mathcal{I} \quad (7)$$

where β^a is the *desired storage profile* given as the optimal storage profile subject to a fixed storage capacity of $\kappa^a(t+1)$ and $\gamma_2 \in (0, 1]$ is the learning rate of the strategy. Note that we analyse in more detail the sensitivity of the learning parameters γ_1 and γ_2 as part of the empirical study of the system in the next section.

6. An Empirical Analysis of Micro-Storage in the UK Market

Based on the adaptive storage strategies defined in the previous section, the analysis of micro-storage we present in this section aims to complement the theoretical part of this paper which assumed a largely homogeneous population of agents (see Section 4). In particular, here we evaluate the emergent properties of a large populations of 1000 agents¹⁵ owning micro-storage of different charging and discharging rates and sizes and using our learning strategy to adapt their storage profile during 500 simulation days over 200 runs. At the beginning of each simulation day, an agent makes a prediction of its load profile and the market prices (using historical data) across the 48 time-slots to compute its best response on a day-ahead basis. At the end of the simulation day, the actual market prices are computed from the total domestic market demand (i.e., the aggregate load and storage profiles) based on the market macro-model described in Subsection 3.2 and published to all the agents *a posteriori*. To frame our results within a real-world context, our simulations focus on the UK electricity grid.¹⁶

Thus, given our macro-model of the UK electricity retail market (see Section 3.2), we initialise individual consumers with typical UK load profiles.¹⁷ Moreover, the learning rates of the agents (presented in Sections 5.1 and 5.2), as well as their charging and discharging rates, are normally distributed around means drawn from charging (and discharging) rates of current technologies (see Section 2).¹⁸ This is done to represent heterogeneous consumers with different types of storage devices and address a more realistic scenario than in our game-theoretic analysis.

Given this setup, for benchmarking purposes we first compute the competitive equilibrium predicted by our game-theoretic framework (which assumes complete information

15. This simulation could be readily extended to hundreds of thousands of agents given the distributed nature of the computation. Now, given that, on average, agents tend to have similar load profiles, we assume here that there are 1000 different clusters of load profiles among the domestic consumers such that a simulation with only 1000 agents (rather than millions) is valid. We have observed from real data that such an assumption about the heterogeneity of the UK domestic population is reasonable.

16. Note that we choose UK as a typical deregulated market. Our approach is nonetheless general enough that our framework can be applied for other markets, including industrial or commercial markets (as opposed to the residential case we consider) as well as micro-grids or the national grid of other countries based on a macro-model of their electricity market. Thus, the results and insights presented in this paper broadly generalise.

17. The profile of an agent is based on a normal distribution with mean ℓ^{mean} which is taken as the UK average profile. It is formally defined as $\ell_i^a \sim N(\ell_i^{mean}, \sigma)$, $\ell_i^a \geq 0 \quad \forall i \in \mathcal{I}$ where $\sigma = 0.2$ to approximate the typical spiky daily profiles of consumers.

18. In all experiments except for when we analyse the effect of learning rates, we draw values for the learning rates from a normal distribution $N(0.05, 0.02)$. Through experimentation, this was found to result in good system-wide performance (see Section 6.5) and is not too small that the system is slow to converge.

about all agents' load profiles and battery capacities) for the UK electricity grid. Second, we study how the system breaks down (in terms of the average individual costs and grid performance metrics defined in Section 3.3) if agents were to gradually adopt storage with no adaptive mechanism. This comparison motivates the need for our adaptive mechanism, and, against these results, we empirically demonstrate that the market does indeed converge to the competitive equilibrium when the population of agents adopt our adaptive storage mechanism. Third, we perform a sensitivity analysis of the convergence properties of our adaptive mechanism with respect to the learning rates to understand the impact of different learning rates on cost savings and grid performance and how these parameters should be set. Fourth, we evaluate the robustness of our approach when agents with micro-storage do not adopt our learning mechanism. Finally, we evaluate the impact of different degrees of saturation of micro-storage (using our agent-based micro-storage management) on the efficiency of the grid and, in so doing, study the performance of the grid (see Subsection 3.3) as more and more consumers adopt the technology.

6.1 The Game-Theoretic Solution

Given the game-theoretic framework outlined in Section 4, we first calculate the competitive equilibrium based on the average domestic consumption profile of a consumer (from the UK) using the procedure described in Subsection 4.2 (i.e., where devices are perfectly efficient and costless). The resulting storage profile is shown in Figure 4. It is clear that the equilibrium behaviour for a consumer is to charge at off-peak hours (at night) and use the stored energy during peak hours (after working hours) when the consumers' load is highest. Furthermore, as observed in Figure 4, at the equilibrium, the optimised load profile (i.e., the sum of the aggregate unoptimised load profile and storage profile) is flattened completely with a load factor of 1. This implies that agent-based micro-storage management can theoretically reduce peaks completely and be completely efficient. Furthermore, the storage capacity required to achieve this equilibrium is 2.3 kWh, computed as the maximum of the cumulative sum of the storage profile (as the micro-storage device charges and discharges over each time-slot). In the next subsection, we first demonstrate how, in practice, the system breaks down completely without an adaptive mechanism to motivate our need for an adaptive mechanism and, we then go on to show that the system indeed converges to the competitive equilibrium when our decentralised adaptive mechanism is adopted.

6.2 Market without the Adaptive Mechanism

To analyse how the market operates without the adaptive mechanism, we set up a population of agents playing their best-response storage profile every day (i.e., using the optimisation algorithm defined in Section 5.1). Moreover, we simulate the gradual adoption of storage devices by the consumers, with a rate of adoption, r (i.e., a probability of r that an agent will adopt storage and keep a storage device; $r = 1$ simulates a system where all agents have storage capabilities at the beginning). Once an agent has storage capability, it will optimise daily and use its best-response storage profile. For this setting, Figure 5(a) shows the deviation from the competitive equilibrium while Figure 5(b) shows the load factor of the grid for different values of r . When $r = 1$ (i.e., all consumers adopt micro-storage





Figure 4: Storage profile and load profile at the competitive equilibrium.


devices at the same time), the system clearly deviates from the equilibrium with the load factor jumping immediately from 0.66 (without micro-storage) to 0.4 (with the immediate adoption of microstorage), suggesting larger peaks in the system. For smaller values of r , the system converges at the beginning (as the demand slowly decreases at peak time and increases at off-peak time since only a small proportion of agents can change their demand), only to break down after a number of simulation days and ends up with larger peaks. The smaller r is, the longer the system takes to break down, though it inevitably does so. The intuition behind this is that there invariably reaches a point when there are too many agents that have adopted micro-storage devices and are using their best-response storage profile. With too many agents re-optimising their storage profiles at the same time, the peaks in market demand are moved around such that aggregate demand profile is not flattened (inferred from the non-increasing load factor on a long-term). In the next subsection, we show that our adaptive strategy helps remedy this and results in desirable system-wide performance.

6.3 Market with the Adaptive Mechanism

Here, we study the convergence properties of our adaptive mechanism and also show how the results corroborate the theoretical bound on storage capacity suggested in Section 4 (given the worst case scenario in Subsection 6.2 when $r = 1$). In more detail, given a population of agents using the adaptive mechanism with $\gamma_1 \sim N(0.05, 0.02)$ and $\gamma_2 \sim N(0.05, 0.02)$ (we show how the performance varies with different learning rates in the next section), the average storage profile is found to converge rapidly to the competitive equilibrium of our



(a) Deviation of the population without adaptive storage from competitive equilibrium.



(b) Load factor (LF) in a market with no adaptive mechanism.

Figure 5: Convergence properties of a system with no adaptive storage.

game-theoretic analysis within less than 20 simulation days (see Figure 6).¹⁹ As expected, the convergence results from the agents gradually adapting their storage profiles such that the aggregate market demand is shifted from peak to off-peak.




Figure 6: Convergence of the average storage profile to the competitive equilibrium when all agents adopt the adaptive mechanism.

As the system converges to the competitive equilibrium, we also observe how the grid efficiency (as measured by the LF and DF) improves and gradually converges (see Figure 7) as agents adapt their storage profiles and the system converges to the competitive equilibrium. In more detail, from Figure 7, we observe that the system LF increases from 0.68 and converges to around 0.93, suggesting considerably fewer peaks in the grid when micro-storage is adopted, and indeed, a flattened demand.²⁰ This is coupled with a DF that is close to 1 which indicates (as discussed in Section 3.3) that, even though agents have closely correlated load profiles, overall, these profiles tend to be reasonably flat.

Furthermore, from Figure 8, we can see that the average storage capacity required converges to around 2.3 kWh (which equals the storage capacity prescribed by our analytical solution — see Section 6.1) after several simulation days. This implies that while the average

19. Given that weekdays are homogeneous (as opposed to Saturdays and Sundays), the agent can learn across weekdays such that the system would converge within a couple of weeks.

20. The results for each simulation day were averaged over the number of simulations. Furthermore, a simulation size of 200 was statistically significant, with results in the figures given with error bars at the 95% confidence interval.




Figure 7: Efficiency of the grid as system converges to competitive equilibrium when all agents adopt the adaptive mechanism.

consumer may buy a storage device of capacity 3kWh (see the maximum storage capacity in Figure 8), the agent would actually only need 2.3kWh of this capacity to minimise costs.

6.4 Sensitivity of the Adaptive Mechanism

One of the assumptions of our approach is that agents are expected to adopt the adaptive mechanism. While this can be imposed as a feature of the smart meter controlling the micro-storage device, it can exceptionally be the case that the smart meter is tampered with or that the user programs it to ignore the learning mechanism and only use its best response, i.e., the agent always executes its optimal behaviour. Thus, in this subsection, we study the effect on the system if part of the population were not to adopt our proposed adaptive mechanism, assuming that the whole population has storage capability.

From Figure 9, we can observe that the system is particularly robust and only starts degrading when more than 60% of the population do not adopt the adaptive mechanism and deliberately execute their best response. As the proportion of the population playing their best response increases, the load factor slightly increases to 0.94 until the proportion of population reaches around 60% after which the load factor rapidly decreases to 0.4 (suggesting large peaks in the system) when *all* agents adopt their best response, i.e., there is no adaptive behaviour in the system. While some agents are using their best response, other agents are gradually adapting their storage profiles (implicitly adapting to the impact of the former agents' best-response profile). The system eventually breaks down when too




Figure 8: Average storage capacity required for a typical consumer as system converges to the competitive equilibrium.

many agents are using their best response and too few their adaptive mechanism. We also notice a small increase in load factor as a result of the increased diversity among agents, an emergent behaviour which we again observe when analysing the load factor of a population with different proportions having storage capability (see Subsection 6.6 for a discussion). Note that while the system remains robust upto 60% population saturation of micro-storage even without a learning mechanism, our mechanism ensures that the system does not break down for any proportion of the population, as we observe in Subsection 6.6.

6.5 Sensitivity of the Convergence Properties

To guarantee the consistency of our results given different parameter settings, in this section we explore how the values of the learning rates²¹ γ_1 and γ_2 affect our convergence results.²² In so doing, we aim to determine how fast an agent should ideally adapt its storage profile to maximise its savings while, if possible, helping to improve the efficiency of the grid. Figure 10(a) shows that the smaller the learning rate, the more efficient the system (with a higher load factor). The intuition behind this is that a small learning rate allows the market prices to change gradually. A higher learning rate, on the other hand, would result

21. Because the learning rates are intrinsic to the adaptive mechanism, we assume that γ_1 and γ_2 share the same value without loss of generality as we are interested in how fast our mechanism adapts rather than specifically in the two-part adaptation.

22. The mean load factors and savings between Day 400 and Day 500 (by which time the system generally has converged) were recorded and averaged over 200 runs for different learning rates.





Figure 9: Grid efficiency when consumer switch to their best-response.

in agents adopting their optimal storage profile too quickly rather than gradually, which clearly results in poor savings and poor system efficiency with peaks cycling in the system. As the learning rate increases, the load factor quickly decreases to 0.59 when all agents adopt their best-response immediately. From an individual perspective, the agent would typically set its learning parameter based on its savings. From Figure 10(b), we can see that, likewise, the smaller the learning parameter is, the better the average savings²³ of the individual agent.


Now, because an infinitely small learning rate is infeasible as it implies an infinitely long time to reach the equilibrium, a trade-off is required. Specifically, because the learning parameters are not very sensitive when they are small, a value of 0.05 is reasonable given that the decrease in savings would be negligible. As mentioned in Footnote 18, we use this value for all our experiments.

Given these results, we can claim that our adaptive strategy sets the benchmark for *any learning strategy* in this system (i.e., the base requirement for any such strategy would be convergence). While these results mainly hold for a whole population owning storage, it is important to see how the system performs as storage is gradually introduced in the population. This should enable us to identify the optimal level of adoption of storage that is required in order to maximise grid efficiency. In so doing, we also aim to complement our analytical results (see Section 4.4) with empirical evidence showing the level of adoption that is still profitable for individual users to acquire storage in the UK market.

23. The average saving of a consumer is computed as the difference in her average costs (after the system has converged) in a system with micro-storage (i.e., after the system has converged) and in one without.



(a)



(b)

Figure 10: The sensitivity of the learning rate, $\gamma_1 = \gamma_2$, on the system in terms of (a) the load factor and (b) savings for agents with storage.

6.6 Grid Efficiency with Incremental Micro-Storage Adoption

In the following set of experiments, we investigate the grid performance metrics and carbon emission reductions (which is one of the main aims of this work) achieved by our mechanism as micro-storage is incrementally adopted.²⁴ From Figure 11, we can observe that the grid efficiency peaks at only 32% of the population (rather than at the aggregate storage capacity of the population at 100%). This suggests that only 32% of the population is required to have storage to maximise the grid efficiency. As that proportion increases from 0, more agents have storage capability and thus storage profiles that are being adapted to flatten the demand. Their aggregate storage profile gradually flattens the aggregate load profile such that the load factor increases as does the diversity factor (since more agents now have a storage profile and, thus, an adapted demand profile). Eventually, the load factor peaks at 32%, at which point, the system is flattened as much as it could be. As more agents acquire storage devices, the diversity factor decreases to 1 as more agents use storage, and finally settles at 1, where, on average, they have the same (flatter) load profile and storage profile. Furthermore, with more agents optimising at the same time, the load factor also decreases slightly as too many agents are now trying to optimise and adapt their storage profiles at the same time. Specifically, agents are optimising a *surplus* of storage capacity that is not required to flatten the demand of the grid. As can be seen, the diversity factor decreases which suggests that optimising the surplus of storage increases correlation among individual demand profiles. The increased correlation further suggests that any small peak in the load profiles has more impact which, on average, decreases the load factor.

From Figure 12, we also observe that a significant benefit of storage at a macro-level (i.e., ignoring the individual benefits of the agents – which we study in the next section) is that if a sufficient proportion (32%) of the population does adopt storage, the carbon emissions of the electricity market would decrease appreciably as peak demands are reduced. Indeed, the carbon emission in the UK can be reduced by up to 7% (from 63 to 58.3 kilotonnes CO₂ per day),²⁵ reaching a minimum when the domestic load factor is maximised (since reducing the peaks in demand has the effect of reducing the carbon intensity of the supplied electricity, which in turn, reduces the total carbon emissions).

7. Cost-Benefit Analysis of Micro-Storage

So far, we have studied the efficiency of the grid achieved by the population as storage is gradually adopted. We now turn to the individual consumer who is principally driven by how much profit she can achieve by adopting micro-storage (at a certain cost). In particular,

24. Note that the results we describe here (i.e., the grid efficiency after the system converges) are similar to our game-theoretic solution given that a system where only a proportion of the population has storage translates to a model in which there are smaller aggregate charging and discharging limits — see Section 4.

25. We calculate the carbon emissions by considering the reduction in carbon intensity of the electricity supply when the domestic load factor reduces (see Figure 3). We consider a total population of 26M UK households and scale these results to take account of the fact that these domestic consumers represent 30% of the total UK demand (as discussed in Section 3) and the remaining 70% consisting of commercial and industrial profiles that remain unchanged.




Figure 11: Grid efficiency (i.e., the load factor and the diversity factor) for different proportions of the population using storage.




Figure 12: Total daily carbon emissions from the UK domestic sector for different proportions of the population using storage.

the question we wish to address is that: at what level of adoption is storage still profitable to agents in the system?²⁶ Hence, in our experiments, we first assume no cost of storage and only vary the proportion of the population with storage and record the cost-savings achieved by the parts of the population that have and have not adopted micro-storage. From our results (see Figure 13), we first observe that when very few agents have micro-storage (i.e. close to 0%), the potential average savings to each agent is close to 14%. As more consumers adopt storage, the average saving gradually decreases to slightly less than 8% (as market prices also flatten and agents can no longer benefit from the difference between low off-peak prices and high peak prices). Interestingly, we also observe that consumers without micro-storage also benefit from its use by other consumers. This is because, as empirically demonstrated in Section 6, the adoption of micro-storage flattens the peaks in demand in the system and, thus, market prices. This means cheaper electricity for the domestic market and, indeed, as more and more consumers adopt micro-storage, the savings of those consumers that do not adopt it, also increase. There reaches a point when 48% of consumers adopt micro-storage and the savings for those consumers with and without the technology are equal. As the percentage increases past 48%, the savings for those consumers who do not adopt micro-storage exceed those of the consumers who do adopt it. This is because the consumers with micro-storage are trying to optimise a surplus of storage as argued in Section 6.6. The implication of this dynamic is that consumers are incentivised to adopt micro-storage until the 48% mark is reached. At this point, there is no incentive for the consumers to deviate from their chosen behaviour (i.e., use micro-storage or not).²⁷ Thus, over a long term, the system will converge to an equilibrium where only 48% of the population adopt micro-storage and at that percentage, consumers can expect a saving of 9% on their individual electricity bills (which equates to an annual saving of £60 per household – based on an average annual electricity bill of £675). This result also points to a slight misalignment between the optimal level of storage for the grid (in terms of grid efficiency) and that for the consumers (in terms of savings). In particular, given the results in the previous section (see Figure 11), we note that the 48% level of adoption equates to a domestic load factor of 0.91, while the maximum load factor achievable (assuming control over the proportion of population adoption storage) is 0.94 which occurs at 32% adoption. This suggests that at the proportion of adoption that the system eventually settle at, the system is only slightly suboptimal.

We next consider this dynamic within a more realistic setting when there is a cost to storage, typically a startup cost from hardware installation (e.g., wiring and converter) and the cost of the actual battery. Based on how much storage capacity the average consumer would require to maximise her savings for different proportions of the population, we calculate the savings minus the cost of storage (based on a typical battery costing £200 per kWh and £600 per kWh respectively with a lifetime of 10 years and a fixed startup cost of £200 for both) and assuming an average cost of electricity of £675 for a consumer per

26. Note that, with respect to an average consumer’s savings, this set of experiments complement those in Section 6.5 where we studied the cost-savings of the users as the learning rate is varied.

27. Note that the consumer is aware of its savings with and without micro-storage which can be computed based on its initial load profile and the demand profile with the storage profile.




Figure 13: Savings without storage and with storage (for the different costs of storage).

year.²⁸ Note that savings without storage do not change as they are independent of the cost of micro-storage. Given this setup, our results (see Figure 13) show that there is a clear first-mover advantage. Thus, a maximum saving can be achieved when the proportion of the population with micro-storage is close to zero (i.e., only a few agents in the population own storage). However, the storage capacity these agents require is relatively high at 4.5kWh, decreasing rapidly to 2.3kWh when all consumers in the population adopt micro-storage (see Figure 14). Moreover, based on the savings with cost of storage and the savings without storage, we observe that the equilibrium moves to 23% for a cost of storage device of £200 per kWh and to 10% for a cost of £600 per kWh (given a startup cost of £200). Hence, by combining the latter results with those from Figure 11, we can infer that the grid efficiency quickly drops as the cost of storage increases (as the level of adoption decreases). Thus, the cost of storage can significantly affect the benefits derived by the users and by the grid as a whole.

The above results, however, consider populations of agents drawn uniformly from the UK average load profile. It could therefore be argued that such results may not apply in circumstances where the population is not uniformly distributed and, in particular, if users consume differing amounts of energy — making savings from storage more viable for those consuming more (as they are able to shift more energy across the day and recover the high startup cost of such a system) than others consuming less. Given this, we next expand our

28. We compare the cost of storage with the daily cost of electricity by calculating the daily cost of owning and running a storage device, i.e., the cost of the storage device and the startup costs that we assume, are spread over the lifetime of the storage device which, in this case, is 10 years. A battery costing £200 per kWh and £600 per kWh, both with a startup costs of £200 equals 1p and 3p per kWh per day respectively, while, on average, the daily cost of electricity is 180p for an 8kWh daily electricity consumption (see Subsection 2.1).




Figure 14: Average storage capacity required for different proportions of the population with storage.

cost-benefit analysis to consider a fundamental distinction between users, namely the low-end users (typically with a yearly consumption²⁹ of 1650 kWh) and high-end users (typically with a yearly consumption of 4950 kWh). In particular, we want to analyse the dynamics of the proportion of low-end and high-end users that will adopt storage (i.e., given a number of agents of each type, what proportion of each type will adopt storage). The aim here is to investigate which type of consumers can benefit the most from our system and, indeed, whether our system can be more efficient. To this end, we adopt an evolutionary game theoretic approach that is suitable to analyse such dynamics, and determine whether the system eventually settles to a stable equilibrium where the behaviours (whether or not to adopt storage) do not change. We next describe our evolutionary game-theoretic framework and, thereon, provide a cost-benefit analysis for low-end and high-end users.

7.1 The Evolutionary Game-Theoretic Model

Here, we formulate the problem as a game where all low-end users adopt the same mixed strategy³⁰ $\pi_r^{LU} \in (0, 1)$ (i.e., the probability that the low-end users have storage capability) and are only motivated by financial gains, while all high-end users adopt a mixed strategy

29. This data is drawn from typical consumption data published by British Gas UK at www.britishgas.co.uk. Furthermore, we assume that both types of consumers have the same normalised daily average load profiles.

30. We assume that agents of the same type share the same mixed strategy given that, on average, they have the same load and storage profiles and thus, the same expected savings.

$\pi_r^{HU} \in (0, 1)$. By analysing how π_r^{LU} and π_r^{HU} evolve as the payoffs change for different proportions of the population of low-end and high-end users with storage, we want to study how the proportion of the population using storage evolves. To this end, we use the classical evolutionary game-theory (EGT) (Weibull, 1995) in which we first compute the heuristic average payoffs of the low-end and high-end (based on simulations) (whether or not using and not using storage) for different mixed strategies π_r^{HU} and π_r^{LU} , given respectively by:

$$u_{LU}(\pi^{LU}, \pi^{HU}) = \sum_{r \in S} u_{LU}(r, \pi^{LU}, \pi^{HU}) \pi_r^{LU}$$

$$u_{HU}(\pi^{LU}, \pi^{HU}) = \sum_{r \in S} u_{HU}(r, \pi^{LU}, \pi^{HU}) \pi_r^{HU}$$

where $u_{LU}(r, \pi^{LU}, \pi^{HU})$ is the payoff of low-end users adopting the pure strategy r given the low-end users' mixed strategy π^{LU} and the high-end users' mixed strategy π^{LU} and $u_{HU}(r, \pi^{LU}, \pi^{HU})$ the corresponding the payoff for high-end users.

We then use these results to calculate the *replicator dynamics*, $\dot{\pi}^{LU}$ and $\dot{\pi}^{HU}$, that describe the dynamics of the population (i.e., how the proportions of low-end and high-end users are evolving), and are given by:

$$\dot{\pi}_r^{LU} = [u_{LU}(r, \pi^{LU}, \pi^{HU}) - u_{LU}(\pi^{LU}, \pi^{HU})] \pi_r^{LU} \quad \forall r \in S$$

$$\dot{\pi}_k^{HU} = [u_{HU}(k, \pi^{LU}, \pi^{HU}) - u_{HU}(\pi^{LU}, \pi^{HU})] \pi_k^{HU} \quad \forall k \in S$$

Finally, we test whether it converges to any Nash equilibria $(\pi_{nash}^{LU}, \pi_{nash}^{HU})$, points where there are no incentives for either low-end or high-end consumers to deviate from.

$$(\pi_{nash}^{LU}, \pi_{nash}^{HU}) = \arg \min_{\pi^{LU}, \pi^{HU}} \sum_{r \in S} \left(\max[u_{LU}(r, \pi^{LU}, \pi^{HU}) - u_{LU}(\pi^{LU}, \pi^{HU}), 0] \right)^2$$

$$+ \sum_{k \in S} \left(\max[u_{HU}(k, \pi^{LU}, \pi^{HU}) - u_{HU}(\pi^{LU}, \pi^{HU}), 0] \right)^2$$

In the next subsection we provide the results of our EGT analysis.

7.2 EGT Analysis of Micro-Storage Adoption with Low- and High-end Users

The results of the EGT analysis are given in Figure 15 for different costs of the storage device (i.e., £0, £200, £400, £500, £1000 per kWh) assuming a typical lifetime of 10 years for a battery and a startup cost of £200. Thus, we can observe that when storage is completely subsidised (i.e., cost of storage is 0), we have a range of Nash equilibria (along the straight line from $(\pi^{LU} = 0.4, \pi^{HU} = 1.0)$ to $(\pi^{LU} = 0.7, \pi^{HU} = 0)$). Furthermore, from Figure 16, we can observe that at these Nash equilibria, the grid efficiency of the system is very high (the load factor is higher than 0.9).

Now, when startup cost increases to £200 per kWh, we now have a single Nash equilibrium at $(\pi_{nash}^{LU} = 0.07, \pi_{nash}^{HU} = 1)$ with a lower load factor of 0.82 (see Figure 16). This implies that the system eventually converges to a Nash equilibrium where *all* high-end consumers adopt storage while only 7% of the low-end users do so. This is because the high-end users overall make more savings (than low-end users) that cover the daily cost of storage and

the high startup costs. However, as the cost of the storage device increases to just £200 per kWh, storage becomes too expensive for all low-end users, with the Nash equilibrium now at $(\pi_{nash}^{LU} = 0, \pi_{nash}^{HU} = 1)$, but is still economically beneficial for all high-end users. With increasing cost, storage gradually becomes too expensive even for the high-end users, as seen from the change of the Nash equilibrium $(\pi_{nash}^{LU}, \pi_{nash}^{HU})$ from $(0, 0.94)$ to $(0, 0.55)$ when the cost of the storage device is £400 per kWh and to $(0, 0.37)$ when the cost of the storage device is £500 per kWh. Eventually, the cost of the storage device is too high even for the high-end users at £1000 per kWh, with the Nash equilibrium now at $(\pi_{nash}^{LU} = 0, \pi_{nash}^{HU} = 0)$. By considering the change in the Nash equilibria as the cost of the storage device increases, we also observe from Figure 16 that the domestic load factor decreases gradually to 0.68 when micro-storage has no impact on the grid, being too expensive to be adopted. From this analysis, we gather that to improve grid efficiency and to maximise the impact of micro-storage on the grid, the cost of storage has to be sufficiently small, and subsidising storage would help improve the efficiency of the grid.

8. Conclusions

In this paper, we set out to explore the theoretical and practical foundations of agent-based micro-storage implementation in the smart grid. To achieve these objectives, we first developed a game-theoretic framework to analyse the strategic choices that agents make in using micro-storage devices in the grid. Our framework allows one to predict the competitive equilibrium of the system, and in particular, specify theoretical bounds on the level of micro-storage adoption and the capacity of micro-storage that will be adopted by a largely homogeneous population.

Building upon the intuitions generated by our theoretical results, we then went on to devise a novel micro-storage strategy that allows an agent to optimise both its storage profile and storage capacity in order to maximise its owner’s savings. Furthermore, we provided an adaptive mechanism based on predicted market prices that allowed the agent to change its strategy in response to changing market prices. Our empirical evaluation of this mechanism on the UK electricity grid was then shown to cause the average storage profile to converge to the theoretical competitive equilibrium. At that point, peak demands are reduced, reducing the requirements for more costly and carbon-intensive generation plants. Moreover, in our analysis of the grid efficiency at this equilibrium we show that, while being stable, it results in reduced costs and carbon emissions. This also shows that the objective of buying storage to save on electricity bills is generally aligned with maximising the grid efficiency (i.e., flattening the peaks in the demand). In particular, we show that, without the burden of cost (e.g., if storage were completely subsidised), the population would adopt storage until an equilibrium with 48% of the population adopting storage is reached and this is achieved with a high level of domestic load factor (i.e. 0.91). Given this, we analysed the system in a more realistic setting and empirically demonstrated that if costs of storage are not sufficiently low, the system will not converge to an equilibrium with a high grid efficiency, and if costs are simply too high, there will be no incentives even for the high-end electricity users to adopt micro-storage.




Figure 15: Evolutionary game-theoretic analysis for different costs of storage. Lines are trajectories representing the evolution of the proportion of low-end and high-end users adopting storage. The black dots are the Nash equilibria.




Figure 16: Load factor for different proportions of low-end and high-end users adopting storage.

In general, our theoretical and practical results provide fertile ground for research into agent-based techniques that might be applied to manage demand in the smart grid. In particular, demand-side management technologies (Hammerstrom et al., 2008), which involve loads (e.g., washing machine or dishwasher) in a user’s home being automatically scheduled to run at certain times, present similar properties to micro-storage in that they allow energy to be moved around from peak to off-peak times in order to flatten demand and reduce costs. Hence, applying a similar framework and strategies to ours, we could expect analogous theoretical results and efficiency gains being predicted for deployments of such technologies. Moreover, our techniques could be used to predict how demand would generally vary across the day once real-time pricing is rolled out and in different regions (populated by different proportions of low-end or high-end users), and this could help better prepare assets (e.g., spinning reserve or strengthening transmission lines).

To generalise our techniques further, we intend to integrate more sophisticated models of the electricity market mechanism into our work in order to better capture the price fluctuations that can occur in real markets. For the theoretical analysis, we will turn to stochastic processes, which are commonly used to model volatility in financial markets. Further, our experiments can be extended by generating prices by drawing samples from suitable distributions, or existing data points, or even by developing our simulations to include market clearing with strategic bidding by energy suppliers and consumers. Accordingly, we also intend to employ better forecasting models of demand and supply (e.g., using Gaussian processes or other regression techniques) to predict prices in our optimisation model. Indeed, so far we have assumed that agents predict prices simply from the previous days’ prices for individual settlement periods. As shown by Wellman, Reeves, Lochner, and Vorobeychik

(2004), agents can perform significantly differently if they adopt different approaches to price prediction and therefore, to improve our empirical analysis, it will be interesting to see how the grid performance and individual agents' profits are affected as different agents adopt different forecasting models. In particular, it will be important to determine how the widespread adoption of micro-storage devices should affect the volatility of the wholesale electricity market.

Furthermore, we intend to explore mechanisms that can ensure convergence towards the equilibrium. In particular, our point of departure will be the theory of strategic behaviour found in the Minority Game (Challet & Zhang, 1997), which shares similarities with our problem, or using a more complex pricing mechanism where the consumers always play their best response (Voice et al., 2011) such that a learning mechanism for the consumer agent would not be required. As part of this work, we also intend to investigate the market efficiency for different proportions of the population adopting micro-storage devices, and whether the decrease in efficiency observed when the market is saturated was a by-product of our adaptive mechanism which could be avoided.

Finally, we intend to consider the grid distribution network. Because peaks are different across different nodes of the electricity network, more storage capacity might be required in some areas than in others. Thus, we will investigate whether our agent-based micro-storage management approach can be used to flatten peaks locally within a node of the electricity network rather than flattening the *aggregate* demand profile of the grid.

Acknowledgments

This paper extends our previous work (Vytelingum, Voice, Ramchurn, Rogers, & Jennings, 2010). It extends the game-theoretic framework to consider levels of micro-storage adoption and expands the empirical evaluation to consider the sensitivity of our convergence properties to the learning rate and more complex agent populations. This work was funded by the iDEaS project (<http://www.ideasproject.info>).

References

- Bathurst, G. N., & Strbac, G. (2003). Value of combining energy storage and wind in short-term energy and balancing markets. *Electric Power Systems Research*, 67(1), 1–8.
- Challet, D., & Zhang, Y. C. (1997). Emergence of cooperation and organization in an evolutionary game. In *Physica A*, Vol. 246(3-4), pp. 407–418.
- Daryanian, B., Bohn, R., & Tabors, R. (1989). Optimal demand-side response to electricity spot prices for storage-type customers. *IEEE Trans. on Power Systems*, 4(3), 897–903.
- DECC (2009). Smarter grids: The opportunity. Tech. rep., Department of Energy and Climate Change (DECC). <http://www.decc.gov.uk>.
- den Bossche, P. V., Vergels, F., Mierlo, J. V., Matheys, J., & Autenboer, W. V. (2006). Subat: An assessment of sustainable battery technology. *Journal of Power Sources*, 162(2), 913 – 919.

- Exarchakos, L., Leach, M., & Exarchakos, G. (2009). Modelling electricity storage systems management under the influence of demand-side management programmes. *International Journal of Energy Research*, 33(1), 62–76.
- Galvin, R., & Yeager, K. (2008). *Perfect Power: How the MicroGrid Revolution Will Unleash Cleaner, Greener, More Abundant Energy*. McGraw-Hill Professional.
- Hammerstrom, D., et al. (2008). Pacific Northwest GridWise Testbed Demonstration Projects; Part I. Olympic Peninsula Project. Tech. rep., PNNL-17167, Pacific Northwest National Laboratory (PNNL), Richland, WA (US).
- Harris, C. (2005). *Electricity Markets: Pricing, Structures, and Economics*. Wiley.
- Holland, A. (2009). Welfare losses in commodity storage games. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems*, pp. 1253–1254, Budapest.
- Houwing, M., Negenborn, R. R., Heijnen, P. W., Schutter, B. D., & Hellendoorn, H. (2007). Least-cost model predictive control of residential energy resources when applying μ chp. In *Power Tech*, pp. 425–430, London, UK.
- Jennings, N. R., Corera, J., Laresgoiti, I., Mamdani, E. H., Perriolat, F., Skarek, P., & Varga, L. Z. (1996). Using ARCHON to develop real-world DAI applications for electricity transportation management and particle accelerator control. *IEEE Expert Systems*, 11(6), 60–88.
- Kirschen, D., & Strbac, G. (2004). *Fundamentals of power system economics*. Wiley.
- Kok, K., & Venekamp, G. (2010). Market-based control in decentralized power systems. In *Proceedings of the First International Workshop on Agent Technology for Energy Systems*.
- Li, H., & Tesfatsion, L. (2009). Development of open source software for power market research: The AMES test bed. *Journal of Energy Markets*, 2(2), 111–128.
- Lund, H., & Kempton, W. (2008). Integration of renewable energy into the transport and electricity sectors through V2G. *Energy Policy*, 36(9), 3578–3587.
- Ramchurn, S., Vytelingum, P., Rogers, A., & Jennings, N. R. (2011a). Agent-based control for decentralised demand side management in the smart grid. In *Proceedings of the Tenth International Conference on Autonomous Agents And Multi-Agent Systems*, pp. 5–12.
- Ramchurn, S., Vytelingum, P., Rogers, A., & Jennings, N. R. (2011b). Agent-based homeostatic control for green energy in the smart grid. *ACM Trans. on Intelligent Systems and Technology*, 2(4).
- Ramchurn, S., Vytelingum, P., Rogers, A., & Jennings, N. R. (2011c). Putting the smarts into the smart grid: a grand challenge for artificial intelligence. In *Communication of the ACM (to appear)*.
- Rogers, A., & Jennings, N. R. (2010). Intelligent agents for the smart grid. *PerAda Magazine*, 10.2417/2201005.003002.
- Scheweppe, F., Tabors, R., Kirtley, J., Outhred, H., Pickel, F., & Cox, A. (1980). Homeostatic utility control. *IEEE Trans. on Power Apparatus and Systems*, 99(3), 1151 – 1163.

- Schwepe, F. (1988). *Spot pricing of electricity*. Kluwer academic publishers.
- Shibata, A., & Sato, K. (1999). Development of vanadium redox flow battery for electricity storage. *Power Engineering Journal*, 13(3), 130–135.
- Smith, K. (2010). Energy Demand Research Project - Fifth Progress Report. Tech. rep., OFGEM, UK. <http://www.ofgem.gov.uk/Pages/MoreInformation.aspx?docid=19&refer=sustainability/edrp>.
- Sovacool, B. K., & Hirsh, R. F. (2009). Beyond batteries: An examination of the benefits and barriers to plug-in hybrid electric vehicles (phevs) and a vehicle-to-grid (v2g) transition. *Energy Policy*, 37(3), 1095–1103.
- Sun, J., & Tesfatsion, L. (2007). DC optimal power flow formulation and solution using QuadProgJ. In ISU Economics Working Paper No. 06014.
- Swider, D. (2007). Compressed air energy storage in an electricity system with significant wind power generation. *IEEE trans. on energy conversion*, 22(1), 95.
- US Department Of Energy (2003). Grid 2030: A National Vision For Electricity's Second 100 Years.. <http://www.oe.energy.gov/smartgrid.htm>.
- van Dam, K. H., Houwing, M., & Bouwmans, I. (2008). Agent-based control of distributed electricity generation with micro combined heat and power—cross-sectoral learning for process and infrastructure engineers. *Computers & Chemical Engineering*, 32(1-2), 205–217.
- Voice, T. D., Vytelingum, P., Ramchurn, S., Rogers, A., & Jennings, N. R. (2011). Decentralised control of micro-storage in the smart grid. In *Proceedings of the 25th International Conference on AI (AAAI)*.
- Vytelingum, P., Voice, T. D., Ramchurn, S., Rogers, A., & Jennings, N. R. (2010). Agent-based micro-storage management for the smart grid. In *Proceedings of the Ninth International Conference on Autonomous Agents And Multi-Agent Systems*, pp. 39–46.
- Weibull, J. W. (1995). *Evolutionary Game Theory*. MIT Press, Cambridge, MA.
- Wellman, M., Reeves, D. M., Lochner, K. M., & Vorobeychik, Y. (2004). Price Prediction in a Trading Agent Competition. *Journal of Artificial Intelligence Research*, 21, 19–36.
- Williams, E. (1984). *DINORWIG: The Electric Mountain*. Central Electricity Generating Board.
- Williams, J., & Wright, B. (1991). *Storage and Commodity Markets*. UIT, Cambridge.
- Ygge, F., Akkermans, J. M., Andersson, A., Krejic, M., & Boertjes, E. (1999). The HOME-BOTS System and Field Test: A Multi-Commodity Market for Predictive Power Load Management. In *Proceedings Fourth International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*, Vol. 1, pp. 363–382.
- Ygge, F., & Akkermans, H. (1999). Decentralized markets versus central control: A comparative study. *Journal of Artificial Intelligence Research*, 11, 301–333.