

RESEARCH

Incentivising Prosocial Behaviour in Community Energy Using Multi-agent Systems

Nathan A. Brooks · Simon T. Powers · James M. Borg

Received: 4 March 2025 / Revised: 30 June 2025 / Accepted: 27 October 2025

© The Author(s) 2025

Abstract

Community energy systems, where communities own their renewable energy sources, are key to the energy transition. But to effectively exploit renewable energy, communities need to reduce their peak consumption. For households, this involves spreading the use of high-power appliances, like washing machines, throughout the day. Traditional approaches rely on differential pricing set by utility companies, but this has been ineffective and raises issues of fairness and transparency. To address this, we investigate a decentralised agent-based mechanism. Agents, representing households, are initially allocated time-slots for when to run their appliances, and can then exchange these with other agents to try and better meet their own preferences. Previous work found this to be an effective approach to reducing peak load when social capital—the tracking of favours—was introduced to incentivise agents to accept exchanges that do not immediately benefit them. We expand this here by implementing appliance usage data from the UK Household Electricity Survey, to determine conditions under which the mechanism can meet the demands of real households. We also demonstrate how smaller and demographically diverse populations of households, with heterogeneity in their demand patterns, can optimise more effectively than larger communities, and discuss the implications of this for designing community energy systems.

Keywords Social capital · Reciprocity · Community energy system · Social learning · Multi-agent systems

1 Introduction

In response to anthropogenic climate change, many countries and international organisations have committed to legally binding greenhouse gas emissions targets. The UK and the EU have both recently updated their legislation to include net zero emissions targets in place for 2050 [1, 2]. This requires moving away from using fossil fuels for energy generation and moving towards renewable sources such as photovoltaic cells and wind turbines. In the UK this is reflected in the proportion of UK energy supplied from low carbon sources rising from 10.1% in 2010 to 20.7% in 2022 [3].

However, increasing reliance on renewable energy raises a key challenge: balancing supply and demand. Traditionally, centralised national grids managed this by adjusting supply - switching fossil fuel power plants on and off to meet users' needs. But as more energy is generated from weather-dependent renewable energy sources,



this approach becomes less viable. How can load balancing be managed when supply is inherently intermittent? This problem is more tractable at smaller scales. As a result, governments and energy providers are increasingly supporting the development of community energy systems – local networks where small towns or neighbourhoods own, manage, and share renewable resources [4, 5]. The European Union’s renewable energy strategy, for example, establishes a legal framework for community energy [6].

Decentralised community energy systems allow for a higher share of renewable technologies to be integrated into energy generation [7], minimise transmission losses between the source of energy generation and the end users [8], and improve energy security as the energy supply is less impacted by geopolitical factors [9]. As social awareness of environmental issues increases, the willingness of communities to invest in community energy systems is also expected to increase [10]. But while there are clear benefits to widespread adoption, the shift towards community energy systems means that communities now become involved in some of the tasks that were previously handled by a centralised national grid. In particular, they now become involved in the balancing of supply and demand from their shared renewable energy sources.

A key problem is how to reduce *peak* demand, i.e. the maximal amount of electricity that is demanded at any one moment in time. If a community’s peak demand is too high then it is unlikely that it will be able to be met by the community’s renewable energy sources, and so the community is likely to have to resort to buying in electricity from non-renewable sources. On the other hand, if the demand could be spread out more evenly throughout the day then all of it may be met from their renewable sources. In the energy literature, this is known as demand-side management – changing user demands to meet the available supply [11]. The opposite approach is supply-side management – changing the amount of energy supplied to meet the demand. However, doing so is often infeasible when weather-dependent renewable energy sources are used. Consequently, mechanisms to try to adapt the energy demands of users to the available supply have come to the fore in recent years, including in UK government policy [12].

In this paper, we examine an agent-based mechanism for demand-side management that does not rely on charging households more to use appliances at certain times, which is often perceived as unfair. We build on preliminary work that proposed a mechanism based on reciprocal exchange of energy usage time-slots between households [13, 14]. However, this work relied on artificially flat demand curves that do not reflect real-world patterns in the times when households use electricity. As such, it has been unclear whether this non-monetary approach to demand-side management can work in real-world conditions.

To resolve this, we address here two research questions: (1) Can a non-monetary, agent-based exchange mechanism effectively perform demand-side management under real-world appliance usage patterns? (2) How does the size and demographic diversity of a community affect performance? To answer these, we incorporate real-world appliance usage data from the UK Household Electricity Survey [15] into our analysis. Our results show the robustness and scalability of the mechanism to real-world community energy systems. They demonstrate that agent-based approaches to demand-side management do not have to rely on changing prices. Instead, we show that a more cooperative approach based on reciprocal exchange of time-slots can be an effective approach for community energy systems.

2 Background

2.1 Pricing-Based Mechanisms for Demand-Side Management

A commonly proposed solution for demand-side management is differential pricing. The idea is that households can be incentivised to run their appliances at times of low demand through lower pricing at these times [16–18]. This has traditionally involved utility companies offering cheaper electricity at fixed off-peak times, such as overnight.

The increasing adoption of smart meters into households has further introduced the possibility of real-time pricing, where prices could vary at each hour of the day based on predicted supply and demand [19]. This idea is at the core of most *agent-based approaches* to demand-side management. With agent-based approaches, households have an intelligent agent program running on a home gateway, which can automatically schedule appliances to run at different times based on price signals (see [20, 21] for reviews).

In theory, a community energy system could use a differential pricing mechanism. However, increasing prices at times of peak demand risks disadvantaging certain types of households, such as those with lower incomes or elderly residents, and magnifying inequalities in access to energy [22]. Charging more to access a resource at peak times is often perceived as unfair by people [23], and people generally rank pricing as one of the least fair resource allocation methods [24].

Moreover, setting higher prices to access the community's renewable energy sources at peak times may lead to people perceiving this as a fine for peak use. Perversely, this can lead to *more* people using the resource at peak times, as they feel justified in doing so by paying for it – a behavioural economics phenomenon known as “a fine is a price” [25–27]. This suggests that using pricing for resource allocation in a community energy system may erode the cooperation and prosocial norms of behaviour that are necessary for the community energy system to work in the long term [10, 28]. Other price-based mechanisms, such as auctions [29], similarly risk setting up competitive rather than cooperative interactions between community members.

2.2 Behavioural Interventions

In contrast to pricing-based mechanisms, an alternative approach from behavioural science uses non-monetary interventions to influence electricity consumption habits [30]. This is often viewed as more socially acceptable than differential pricing. Examples include providing households with *individual feedback* on their energy consumption, providing them with *social information* on their energy use compared to similar households, and gamification techniques such as setting individual and neighbourhood goals to reduce energy consumption [31–33]. Field experiments have demonstrated that these non-monetary incentives can be an effective approach to demand-side management [34, 35].

However, two issues arise when applying these techniques to community energy systems. First, non-monetary incentives need to be done in a way that is sensitive to the social psychology of the households being targeted. Behavioural science research shows that people compare themselves to how well other people are doing when these kinds of interventions are used, and that perceived inequality (“why should I change my behaviour if you don't change yours”) can be a behavioural barrier to changing energy usage [36]. Second, behavioural interventions are limited in their flexibility, typically requiring households to manually change their electricity usage on each day. In contrast, agent-based systems can automate scheduling of appliances, reducing the burden on households.

These limitations point to an opportunity: to combine the social sensitivity of non-monetary incentives [37] with the flexibility and automation of agent-based appliance scheduling. Such a hybrid approach could facilitate fair demand-side management without imposing heavy burdens or monetary penalties on households.

2.3 Agent-Based Appliance Scheduling with Non-monetary Incentives

To try and achieve fair and cooperative demand-side management, Petruzzi et al. [13] and Brooks et al. [14] developed a mechanism to implement non-monetary agent-based appliance scheduling in a community energy system. They considered a scenario where a number of households share a renewable energy source [38]. This scenario is important because sharing photovoltaic cells (for example, among apartments in a building) can lower energy costs more than each household having its own private system, and this cost-saving effect grows as the energy community increases in size [39–41]. However, sharing a renewable energy source between households

introduces a resource allocation problem – which households should get access to the community’s renewable energy at a particular time?

To address this, Petruzzi et al. [13] and Brooks et al. [14] assumed that each household has an agent program running on a home gateway connected to their smart meter, into which they can input their preferred time-slots for when they would like to run high-powered but time-flexible appliances, such as washing machines, dishwashers and clothes dryers. The aim is then to allocate actual time-slots to each household agent for when they run their appliances [42].

Petruzzi et al. [13] proposed an allocation mechanism based on the exchange of time-slots between agents. Agents begin with randomly assigned time-slots for when they can power their appliances from the community energy system’s shared renewable energy source. They are then allowed to negotiate and propose time-slot exchanges with other agents, in order to try to get time-slots that better match their own appliance usage preferences. Petruzzi et al. [13] showed that the mechanism was effective at demand-side management when agents used reciprocal exchange through the tracking of favours. That is, when agents accepted time-slot exchange requests from other agents that did not immediately benefit themselves, but that made the other agent more likely to accept one of their own exchange requests in future. They considered this building of favours as an electronic form of social capital amongst agents [43, 44].

Brooks et al. [14] extended this work by addressing the question of whether the mechanism could still be effective when agents could choose to be *selfish*. Selfish agents only accept exchange requests that immediately benefit themselves, rather than engaging in reciprocal exchange. Brooks et al. [14] showed that the mechanism was resistant to exploitation by selfish agents, provided that agents could keep track of who had provided them with a favour in the past, and could preferentially exchange with them. In a practical implementation this result is important, because it means that each household can be free to choose the initial strategy of its agent from the point of view of self-interest, but that even selfish agents will be incentivised to become cooperative in order to improve their performance. This can allow explanations to be given to households in terms of how their agent is acting in their best interest.

Importantly, the exchange mechanism is decentralised. It would also be possible to use a centralised mechanism, where all agents send their time-slot preferences to a central authority who performs the allocation. However, a centralised mechanism would require all agents to truthfully reveal their preferences, which requires trust in the central authority and could compromise privacy [45]. Such a centralised mechanism may also suffer from scalability problems as the number of agents increases. In addition, we suggest that a centralised mechanism is likely to be perceived as less fair by households, since it may be less transparent and intuitive than the simple exchange of favours. In particular, any allocation mechanism, centralised or decentralised, needs to deal with the problem that there are multiple different allocations that will give the same mean satisfaction of agent preferences. The mechanism then needs to be able to choose one of these allocations, and for this decision to be justifiable to the households affected by it. For example, households may ask questions such as “why did another household have all of their preferences met today, whereas only half of mine were?”. Reciprocal exchange – the trading of favours between agents over time – allows explanations to be given based on the reciprocity principle [46]. For example, “You helped the other household today by giving them a time-slot that you initially had but did not plan to use. This means that household will repay the favour to you on a future day, helping your preferences to be satisfied then”.

2.4 Research Gap

The work of Petruzzi et al. [13] and Brooks et al. [14] left unanswered two important questions about the robustness of the non-monetary agent-based mechanism and its adaptability to real-world conditions. First, their models assumed a uniform distribution of time-slot preferences, which fails to reflect real-world patterns of appliance usage. In reality, certain periods – such as early evening – are more likely to experience high demand, with many agents preferring these time-slots. We hypothesise that this will reduce the effectiveness of the exchange

mechanism in satisfying agent preferences, because when there is less complementarity in preferences then there are less opportunities for exchanges that benefit one or both agents. To investigate this, we analyse the mechanism by building a simulation that uses real-world appliance usage data from the UK Household Electricity Survey (HES) [15].

Second, scalability remains an open question. Previous analyses only considered a fixed-size community energy system. When designing a community energy system, it may be possible to engineer the population size, by breaking a large population into smaller populations within which exchanges occur. It is therefore important to determine which population sizes facilitate reciprocal exchanges and thus maximise mean agent satisfaction. We hypothesise that increasing the number of agents will have two effects. First, theory from economics and social science suggests that smaller populations of agents should be more conducive to reciprocal exchange [47]. Second, in larger populations agents may be more likely to find another agent with a time-slot that they need. To investigate this trade-off, we allow the population size in our simulation to be varied, and discuss the implications of varying population size for a real-world implementation of the mechanism.

We also hypothesise that scalability will be affected by the demographic composition of the population. Different demographics can have very different usage patterns for their appliances. For example, a large family will need to use their washing machine more frequently and at different times of day than a young individual who lives alone. A more demographically diverse population would therefore be expected to have a more diverse set of time-slot preferences, creating more opportunities for exchanges and greater mean satisfaction. The HES contains appliance usage data from households across a range of demographics (e.g. working age, pensioner, single-person household, multi-person household). We use this to investigate how effectively the mechanism can adapt to the demand profiles of different demographics, and to determine whether it is more effective at satisfying household preferences when communities are more demographically diverse.


3 Methods

In order to analyse the mechanism experimentally, we have developed the Energy Exchange Simulation (available to download from the following GitHub repository: <https://github.com/NathanABrooks/EnergyExchangeArena>).

3.1 The Energy Exchange Simulation

The algorithm implemented in the simulation is shown in Algorithm 1, and a flowchart representation is provided in Fig. 1.

Fig. 1 Flow chart showing each stage of a run within the *Energy Exchange Simulation*



```

1: simulation ← current simulation
2: d ← current day
3: A ← set of a agents
4: f ← single population counter = 0
5: while simulation.is_complete() == false do
6:   for each a ∈ A do
7:     a.receive_random_allocation()
8:   end for
9:   e ← inactive exchange count = 0

10: while d.is_complete() == false do
11:   V ← set of v adverts
12:   for each a ∈ A do
13:     v ← a.select_unwanted_time_slots()
14:     V.list_advert(v)
15:   end for

16:   for each a ∈ A do
17:     if a.received_request() == true then
18:       go to next agent
19:     end if
20:     if a.satisfaction() == 1 then
21:       go to next agent
22:     else
23:       r ← a.identify_exchange(V)
24:       a.request_exchange(r)
25:     end if
26:   end for

27:   for each a ∈ A do
28:     if a.received_request() == true then
29:       a.accept_exchange_if_approved()
30:     end if
31:   end for

32:   for each a ∈ A do
33:     if a.made_exchange() and
34:       a.agent_type == Social then
35:       a.update_social_capital()
36:     end if
37:   end for

38:   d.no_trades() = true
39:   for each a ∈ A do
40:     if a.made_exchange() == true then
41:       d.no_trades() = false
42:     end if
43:   end for
44:   if d.no_trades() == false then
45:     e = e + 1
46:   else
47:     e = 0
48:   end if
49:   if e == 10 then
50:     d.is_complete() = true
51:   end if
52: end while

53: for each a ∈ A do
54:   a2 ← random agent to observe
55:   if a.satisfaction() < a2.satisfaction()
56:     then
57:       x ← random value between 0 and 1
58:       if a.learning_probability(a.satisfaction())
59:         > x then
60:         a.copy_strategy(a2)
61:       end if
62:     end if
63:   end for

64: if s.single_type_remaining() == true then
65:   f = f + 1
66: end if
67: if f == 100 then
68:   simulation.is_complete() = true
69: end if
70: d = d + 1
71: end while

```

Algorithm 1 Pseudo-code for the *Energy Exchange Simulation*.

The simulation iterates for a number of days, as follows. Each day, every agent *i* requests four hour-long time-slots in which they require electricity. This defines their set of time-slot *preferences*, denoted by Φ_i , where $|\Phi_i| = 4$. These preferences are generated from *demand curves*, as described in Sect. 3.2.

The aim of the algorithm is to allocate a set of actual time-slots T_i to each agent *i*, such that the difference between the actual and preferred time-slots is minimised. We define the *satisfaction* of agent *i*, s_i , as the proportion of its time-slot preferences that have been satisfied:

$$s_i = \frac{|\Phi_i \cap T_i|}{|\Phi_i|}. \quad (1)$$

We track the mean value of s_i , which we refer to as *mean satisfaction*, as a measure of how well the mechanism is satisfying the agents' preferences. We can calculate a theoretical optimum for the mean satisfaction on a particular day, by comparing all of the requested time-slots at the start of the day with all those available within the simulation. This allows us to compute the maximum % of the agents' time-slot preferences that can possibly be satisfied, and so provides an upper bound on the performance of a time-slot allocation mechanism.

Allocation of time-slots into T_i proceeds as follows. Time-slots are initially allocated to each agent at random at the start of the day, with each agent given four time-slots drawn from a uniform distribution (line 7 in Algorithm 1). This randomness in the initial allocation allows for privacy, since preferences do not have to be shared with the agent doing the initial allocation. It also produces maximal reduction in peak consumption, by spreading consumption evenly across the time-slots. However, agents are highly unlikely to have this random allocation match all of their requested time-slots. To address this problem, after the initial allocation agents can take part in pairwise exchanges where one agent requests to swap one of its time-slots with a second agent, and the second agent decides whether or not to fulfil the request.

Exchanges begin every day once each agent has received their initial allocation and decided which of these time-slots they wish to keep. They then anonymously advertise slots that they have been allocated but do not want to an ‘advertising board’ (lines 11–15 of Algorithm 1). Several exchange rounds then take place during the day, with exchanges continuing until no request has been accepted in the past 10 rounds of exchanges (line 48) and the mean satisfaction has therefore stopped increasing. In each exchange round, agents can request a time-slot from the board so long as they have not already received a request from another agent during that round (lines 16–26). Agents accept or refuse requests (line 29) based on their *strategy*.

Agents can follow either a *social* or a *selfish* strategy when deciding whether to accept exchange requests from other agents. Selfish agents will only accept exchanges that are in their immediate interest. This means that selfish agents need to be offered a time-slot that they have initially requested (an element of $|\Phi_i|$) and do not already have (not currently an element of T_i) in order for them to agree to the exchange. Social agents also agree to these mutually beneficial exchanges. In addition, social agents also make decisions based on *social capital*, in the form of repaying previous favours given to them by other agents. Specifically, when a social agent’s own exchange request is accepted, they record it as a favour given to them. When a social agent receives a request from another agent who previously gave them a favour they will accept the request, if it does not cause them to lose one of their preferred time-slots (i.e. they do not lose a time-slot in Φ_i as a result of the exchange), and then record that the favour has been repaid (line 34). Favours that have been repaid can not be used again – this social capital is considered spent, and so agents need to continue to offer beneficial exchanges in order to build up social capital. This leads to a system of social agents earning and repaying favours among one another, increasing the number of accepted exchange requests.

Agents can change their strategy via *social learning* (note that both selfish and social agents undergo ‘social’ learning). This occurs at the end of each day after the exchange rounds are complete and works as follows (lines 52–60). Each agent i observes a randomly selected second agent j . If the observed agent j has a higher satisfaction (Eq. 1) than agent i , then i switches to j ’s strategy with a probability:

$$p = \max \left(0, 2 \left(\frac{1}{1 + \exp(-\beta(s_j - s_i))} \right) - 1 \right). \quad (2)$$

This is a truncated, scaled sigmoid, which ensures that an agent will only change strategy if the agent it observes has a higher satisfaction than itself. This is in contrast to the standard Fermi function often used in social learning, where agents can copy a strategy giving worse observed outcomes than their own with a non-zero probability [48]. While this allows for elegant mathematical analysis of long-run behaviour at the population level, our focus is on the immediate behaviour of self-interested agents in real-world settings, who would likely never copy strategies observed to be objectively doing worse than their own.

Equation 2 results in payoff-biased imitation dynamics in which strategies giving higher individual satisfaction are more likely to spread in the population. The value of β determines the strength of selection. With larger β values, smaller differences in satisfaction are more likely to cause agents to change their strategy compared to when β is smaller. This thus allows us to control how sensitive agents are to the difference between their own and another agent’s satisfaction when deciding whether to change strategy. Agents that change from the social to the selfish strategy retain their accumulated social capital.

The day ends after social learning. Only unspent social capital, i.e. social agents' memory of favours, remains between days. The simulation continues until the entire population of agents has adopted a single strategy, which always occurs because the population is finite and there is no mutation in the social learning equation. Once a single strategy has fixed in the population at 100% frequency, the simulation continues for a further 100 days (lines 61–66), so that performance can be evaluated once only one strategy has been present for a prolonged period of time. This analysis is necessary because the long-run performance of the social strategy is dependent on agents having the opportunity to build social capital, which may not occur if the simulation terminates after a small number of days.

3.2 Demand Curves


We used data from the HES to integrate realistic energy demand into the simulation [15]. The HES was conducted between 2010 and 2011, and was then the most detailed monitoring of electricity use ever carried out in the UK, providing real demand curves for specific household appliances from 250 UK households [49]. The survey provides dis-aggregated data about specific types of appliances, and can be sorted by key demographic metrics such as the size of the household, the number of residents, and the residents' ages.

We used demand curves for “switchable” appliances, which refers to appliances that are not inherently time bound and so can potentially be shifted, such as dishwashers, washing machines and clothes dryers [49]. The data for the appliances was taken from the “HES 24-Hour Chooser” tool provided with the dataset, with switchable appliances considered to be those classified under “Washer/Drying”.

The resulting demand curves used in this paper are shown in Fig. 2. We normalised the demand curves so that the peak demand for each curve is scaled to one, allowing for easy comparison between the curves given the assumption that each household requires the same number of time-slots. Agent preferences on each day, Φ_i , were generated by selecting four time-slots for each agent from the demand curve by roulette wheel sampling.

For our base analysis, we use the “All Households” curve, which is the mean demand curve for “Washer/Drying” for all households in the HES. To answer our second research question, we also consider the curves for “Single Pensioners” and “Single Non-Pensioners” taken from the survey. These two groups have clear differences

Fig. 2 Demand curves computed from appliance usage data in the UK Household Electricity Survey [15]. These were adapted for the hourly time-slot model by combining data from each 10 min period into the corresponding hour's time-slot. The demand curves were then normalised so that the peak demand of each curve is one (see text). The demand curves represent usage data taken from “Switchable appliances” such as washing machines, dishwashers, and clothes dryers that can be scheduled and do not typically have to be used at a specific time. “Single pensioners” and “Single non-pensioners” represent demographics of individuals living alone



in their usage patterns, with pensioners having a large spike in usage between 10:00 and 11:00 with usage steadily declining until midnight, and non-pensioners having comparatively more consistent usage throughout the day. A larger number of single non-pensioners used their appliances at night, potentially representing appliances running at a scheduled time or being started before the user went to bed. We therefore selected these two demographic groups for analysis as they had markedly different demand curves, so their pairing had the largest potential to increase agent satisfaction. These two demographics were also approximately the same size in the HES, with the single pensioners group representing 34 households and the single non-pensioner group representing 35 households.

3.3 Parameter Values

Unless otherwise specified, parameters were kept consistent with [13] to allow for a comparison of results. The baseline consists of a population of 96 individual agents, each requesting 1kWh units of energy during four time-slots, with 16 kWh energy available to be allocated in each time-slot. In order to vary the population size, we allow the total availability of energy to linearly scale with the total demand (this assumption corresponds to a community energy system with more members installing proportionately more renewable energy). The default value for the strength of selection, β , was 1. We initialised the simulations with 50% of the population using each strategy. All of the parameter values can be set in the `config.properties` file in the simulation code.

3.4 Statistical Analysis

Because the simulation is stochastic, we repeated each experiment for 100 simulation runs. We compared the performance of agents using the selfish and social strategies, in terms of their mean agent satisfactions, as well as reporting the number of runs in which either the selfish or the social strategy went to fixation in the population (termed *selfish takeover* and *social takeover*, respectively). Where we report statistical significance, we first checked for normality using a Shapiro-Wilk test. This showed the distribution of agent satisfactions to not deviate significantly from the normal distribution ($p > 0.05$) in these cases. We also used Levene's test and found no statistically significant difference in the variances of the selfish and social agent satisfaction distributions ($p > 0.05$) in these cases. Consequently, we tested whether the difference in mean agent satisfaction between selfish and social populations was statistically significant using the independent Student's t-test. We considered the difference to be significant when $p < 0.01$.

4 Results

4.1 Logical Comparison of Strategies

Prior to any experimentation, we can logically deduce that agents choosing the social strategy should increase the mean satisfaction of the agents as a collective more so than agents choosing the selfish strategy. This is because given that no agent gives away a slot that it has requested, and agents only request time-slots that they want, any exchange must increase the satisfaction of the agent making the request, and either have no effect on or increase the satisfaction of the agent receiving the request. With exchanges only having a positive or neutral impact on the agents involved, more exchanges must lead to a greater mean satisfaction. In pure populations, where all agents are utilising the same strategy, strategies that allow for more exchanges will have a higher mean satisfaction. As agents using the social strategy will accept any exchange that a selfish agent would accept but will also make exchanges based on social capital, a population of purely social agents will be more satisfied than a population of purely selfish agents.

In order to confirm this behaviour we ran the simulation with pure populations and compared the mean satisfaction of the agents. After a single day, purely selfish populations finished with a mean satisfaction of 0.683, where a mean satisfaction of 1 would mean that all agents had all their preferred time-slots. Purely social populations performed similarly after a single day with a mean satisfaction of 0.696. This is because after a single day very little social capital will have been accumulated by the agents, and so the number of exchanges will be similar to the selfish populations. After 100 days, once social capital has been built up by the social agents, selfish agents see no notable change in performance while social agents' mean satisfaction jumps to 0.834. The difference in mean satisfaction between pure populations was statistically significant, and the highest performing selfish population had a lower mean satisfaction than the worst performing social population. This confirms our hypothesis that purely social populations will be able to outperform purely selfish populations. In this experiment, the theoretical optimum mean satisfaction (as defined in Sect. 3.1) across the 100 runs was 0.852, which the population of social agents was close to obtaining with their mean satisfaction of 0.834.

4.2 Flat Demand Curve

In order to understand how using real-world demand data impacts the behaviour of the system, we first ran the simulation with a flat demand curve, where agents had an equal likelihood of requesting any hour of the day when choosing their four preferred time-slots. All 100 runs finished with the entire population adopting the social strategy from a starting point of half the agents using the social strategy and half using the selfish strategy. The social takeover – the day in which all agents adopted the social strategy – took a mean of 165 days to occur, with the mean satisfaction of the agents being 0.834 on the day that the selfish strategy was eliminated, and 0.843 when the simulation finished. All agents had been able to build up social capital through using the social strategy 100 days after the social takeover.

4.3 Switchable Appliance Demand Curves from the UK Household Electricity Survey

We next considered the “All Households” demand curve for switchable appliances. With this change, social agents were only able to take over the population in 56 of the 100 runs of the simulation, with the other 44 runs resulting in a purely selfish population. In runs where the population fully adopted the social strategy, the mean time for the social strategy to takeover was 310 days, with a mean satisfaction of 0.707 when the strategy was fully adopted, and 0.716 when the simulation ended 100 days later. The optimum satisfaction was a mean of 0.721 across all the days in all simulation runs, and so when the social strategy does fully takeover the population then the agents are able to perform at a near optimum level. In the 44 runs where the selfish strategy took over the population the mean satisfaction of the agents was 0.578 at the end of the simulation, a clear drop in performance compared to the social strategy, which was found to be statistically significant, with no overlap between the two sets of results. This supports our hypothesis that real-world demand curves decrease the effectiveness of the mechanism at satisfying agent preferences, by increasing the ability of selfish agents to take over. When more agents have similar preferences, there are less opportunities for exchanges, and hence less opportunities for social agents to exchange favours.

The time for one strategy to fix in the population, and hence for the simulation to end, is different on different runs. While our focus is on the equilibrium end states, we can also examine what happens to the satisfaction of agents as either the selfish or social strategy starts to take over. Figure 3 shows a representative run where the selfish strategy took over, and a representative run where the social strategy took over. The specific runs shown represent the median run when the runs that ended in the strategy fixing are ordered by how many days it took for fixation (recall that we always run the simulation until either all agents are social or all agents are selfish). The mean satisfaction values shown on the graphs represent a five-day moving mean in order to minimise the effect of day-to-day variance. In both the run that resulted in a fully social population, and the run that resulted in a fully selfish population, the mean satisfaction was very similar for the first 100 days, with selfish agents performing

Fig. 3 Representative runs of the simulation that resulted in the population becoming either fully selfish (top) or social (bottom). The representative runs are chosen as the median run when all runs that became either fully social or fully selfish are ordered by the time it took to reach that state. The values for each day are calculated as a five point moving mean. The “Optimal Satisfaction” curve gives an upper bound on mean agent satisfaction (see Sect. 3.1). The “Random Satisfaction” curve shows mean agent satisfaction from the random initial time-slot allocations at the start of the day, before exchanges take place




Table 1 Statistics for specific days from the typical social run displayed in Fig. 3. Social Population shows the number of the 96 agents that are using the social strategy. SD refers to standard deviation from the mean satisfaction

Day	1	50	100	150	200	250	300	350
Social population	48	52	66	47	77	92	96	96
Mean social satisfaction	0.641	0.654	0.700	0.654	0.714	0.687	0.716	0.734
Mean social SD	0.244	0.215	0.236	0.239	0.257	0.281	0.257	0.262
Mean selfish satisfaction	0.526	0.590	0.742	0.689	0.711	0.4	0	0
Mean selfish SD	0.193	0.233	0.199	0.217	0.233	0.122	0	0

very similarly to their social counterparts. Later in the simulation however, the selfish run’s mean satisfaction dropped. This is because there were no longer enough social agents in the population to increase the mean satisfaction by performing more exchanges.

Tables 1 and 2 show specific days in the typical simulation runs displayed in Fig. 3. Upon comparing the outcomes after a single day within the simulation, it becomes evident that the two runs exhibit distinct disparities in mean social satisfaction. This observation underscores the considerable variation in the distribution of time-slots between simulation runs, which can be attributed to the initial random assignment of these time-slots. It is also clear that the strategy used by agents fluctuates greatly throughout a typical simulation, with neither run having a steady trend in strategy used towards its final state. This fluctuation greatly benefits the selfish strategy, as

Table 2 Statistics for specific days from the typical selfish run displayed in Fig. 3. Selfish Population shows the number of the 96 agents that are using the Selfish strategy. SD refers to standard deviation from the mean satisfaction

Day	1	50	100	150	200	250	300	350
Selfish population	48	60	40	65	84	88	92	96
Mean selfish satisfaction	0.526	0.629	0.663	0.6	0.628	0.545	0.568	0.596
Mean selfish SD	0.246	0.216	0.241	0.227	0.227	0.243	0.236	0.226
Mean social satisfaction	0.589	0.590	0.665	0.710	0.583	0.563	0.689	0
Mean social SD	0.213	0.237	0.223	0.212	0.236	0.207	0.108	0

Table 3 How population size affects the performance of social and selfish strategies

Population size	24	48	72	96	120	144	168	192
Social takeovers	84	87	82	56	31	11	5	1
Mean takeover days (social)	111	184	229	310	318	367	477	380
Mean takeover satisfaction (social)	0.598	0.663	0.694	0.707	0.722	0.736	0.725	0.731
Mean satisfaction 100 days after takeover(social)	0.623	0.688	0.704	0.716	0.727	0.733	0.738	0.730
Selfish takeovers	16	13	18	44	69	89	95	99
Mean takeover days (selfish)	131	219	277	365	355	334	317	302
Mean takeover satisfaction (selfish)	0.421	0.514	0.555	0.584	0.604	0.617	0.625	0.632

agents can accrue social capital while being social before switching to the selfish strategy and benefiting from their existing social capital with social agents, while no longer assisting others by accepting requests based on social capital. These tables also demonstrate how on a given day there is often little difference between the mean performance of the two strategies, with the better performing strategy often fluctuating day to day. This promotes the fluctuation between strategies, benefiting the selfish strategy as discussed. As we have previously shown, the population fully adopting the social strategy is in the best interest of all agents involved, as it increases the mean satisfaction of the agents. These results show that in order for a population to consistently adopt the social strategy, there needs to be enough of a consistent performance difference between the two strategies that agents are both more likely to adopt the social strategy, and unlikely to switch when observing a selfish agent.

4.4 Effect of Population Size

We investigated the optimum population size for the mechanism by re-running the simulation with a variety of population sizes between 24 and 192 agents. These numbers were chosen in order to be between 25% and 200% of the initial population of 96 agents. These populations continued to use the “All Households” demand curve. Table 3 shows how the system’s behaviour changed by adjusting only the population size. We show mean satisfaction at the point one strategy takes over in the population, and for social takeovers, mean satisfaction a further 100 days after this (to show the effect of pure social populations accumulating social capital). The smallest population of 24 agents adopted the social strategy far more frequently than the larger populations, with the social strategy fixing in 84 of the 100 simulation runs. This is a clear contrast with the largest population of 192 agents, which adopted the social strategy only once in all 100 simulation runs. However, while the smallest population adopted the social strategy more often, the mean satisfaction at the end of the simulation was only 0.623, which is similar to the mean satisfaction seen with the 192 population size selfish runs (0.632).

This result supports the hypothesis that a trade-off must be made when deciding on the population size for a community energy system using this energy exchange mechanism. Smaller populations are more likely to adopt the social strategy. Larger populations are less likely to adopt the social strategy, but their potential performance is higher when they are able to do so. The reason we see this behaviour is because in smaller populations, each individual is more likely to interact with each other individual and so social capital can be built up much more quickly. At the end of the simulation runs, with the smallest population size agents had a mean of 32.55 unspent social capital, whereas in the largest population the mean was only 21.44 unspent social capital. In larger

populations social capital is slower to build, but with more agents there are more time-slots in the system, and so more agents to potentially trade with. It is also worth noting that there was little difference in the number of populations taken over by the social strategy between the simulations with populations of 24, 48 and 72 agents, which were taken over by the social strategy 84, 87 and 82 times, respectively. This shows how decreasing population size only benefits the social strategy up to a certain point, and so it is important not to reduce the population size beyond this in order to maximise mean agent satisfaction, which was higher with larger social populations.

4.5 Effect of Diverse Demographics

We next explored how having mixed demographics within the population influenced performance, by using the demand curves for single pensioners and single non-pensioners (Fig. 2). We ran the simulation such that half the agents would use the demand curve of single pensioners for their requests for time-slots and the other half used the demand curve of single non-pensioners.

In this mixed demographic scenario, the social strategy was able to take over the population in 90 of the 100 simulation runs for a population size of 96. This took a mean of 272 days and resulted in a mean satisfaction of 0.733 immediately after all agents became social, and a mean satisfaction of 0.750 when the simulation ended 100 days later. As the social strategy is able to effectively allocate time-slots for maximum agent satisfaction, mean satisfaction of social agents was typically within 0.01 of the calculated optimum value for any given run. For the 10 simulation runs that adopted the selfish strategy, it took a mean of 385 days for the selfish strategy to take over, and the mean satisfaction was 0.600 at the end of the simulation. The lowest mean satisfaction from a social takeover at the end of a run was higher than the highest mean satisfaction from a selfish run, with the difference between the two types of runs being statistically significant. Comparing these social populations with social populations using a single real-world demand curve (see the 96 population size shown in Table 3) shows a statistically significant increase in mean satisfaction.

By diversifying the population such that they have varied demand curves, the social strategy was able to take over more consistently and the mean satisfaction increased compared to when all agents used a single demand curve. This supports the hypothesis that a demographically diverse population, with a greater heterogeneity of time-slot preferences, increases the effectiveness of the mechanism.

We also considered the effect of population size on this mixed population of 50% single pensioners demand curve and 50% single non-pensioners demand curve. To do this we re-ran the simulation with varying population sizes, as with Table 3, to see whether alternative population sizes notably affect the results. As seen in Table 4, a mixed population allowed for improved mean satisfaction for runs where the social strategy took over, and also allowed the social strategy to take over more consistently compared to when a single demand curve was used as seen in Table 3. This impact was most evident with the larger population sizes -- the populations of 96, 120 and 144 agents improved on the number of runs adopting the social strategy from 56, 31 and 11 runs to 90, 81 and 80 runs, respectively. This clearly shows how a population with diverse demands can perform well with a much larger range of population sizes than a population with more similar demands. It is also worth noting how the smallest population tested, 24 agents, adopted the social strategy less frequently than the slightly larger sizes

Table 4 The effect of population size in mixed populations of single pensioners and single non-pensioners (see text)

Population size	24	48	72	96	120	144	168	192
Social takeovers	85	94	93	90	81	80	57	42
Mean takeover days (social)	107	169	192	272	321	382	409	463
Mean takeover satisfaction (social)	0.609	0.683	0.712	0.733	0.747	0.756	0.758	0.770
Mean satisfaction 100 days after takeover(social)	0.656	0.706	0.731	0.750	0.754	0.758	0.768	0.767
Selfish takeovers	15	6	7	10	19	20	43	58
Mean takeover days (selfish)	108	174	282	385	340	549	552	533
Mean takeover satisfaction (selfish)	0.413	0.505	0.574	0.604	0.617	0.626	0.648	0.654

Table 5 The effect of the strength of selection in mixed populations of single pensioners and single non-pensioners (see text)

β value	0.5	0.75	1	2	3	4	5
Social takeovers	100	95	90	51	38	30	25
Mean takeover days (social)	325	282	272	219	176	152	111
Mean takeover satisfaction (social)	0.737	0.734	0.733	0.735	0.726	0.730	0.726
Mean satisfaction 100 days after takeover(social)	0.741	0.743	0.750	0.749	0.744	0.734	0.742
Selfish takeovers	0	5	10	49	62	70	75
Mean takeover days (selfish)	–	332	385	258	176	145	128
Mean takeover satisfaction (selfish)	–	0.630	0.604	0.599	0.607	0.603	0.606

of 48, 72 and 96 agents. This further demonstrates how it is important to select a population size that allows for populations to adopt the social strategy while being large enough for a high mean satisfaction, but goes beyond our previous understanding by suggesting that having too small of a population size can reduce the social strategies ability to take over consistently. This is because while smaller populations build up social capital more easily, with less potential trade partners there is a greater risk of only a small number of trades happening in a day, and those trades being between selfish agents.

4.6 Effect of the Strength of Selection During Social Learning

Finally, we considered the effect of the agents' learning process by varying the strength of selection, β , in Eq. 2. We ran a population of 96 agents using the two separate demand curves for single pensioners and single non-pensioners from the previous experiment. Table 5 shows how the dynamics and outcomes of the mechanism are affected as β is varied. The results show how as the strength of selection increases, the number of simulation runs in which the social strategy successfully takes over the population decreases. This makes logical sense, as with a greater strength of selection agents will change their strategy more frequently. This will lead to more selfish agents benefiting from social capital they gained from when they were previously social, preventing social agents from having an advantage. We also see that with a greater strength of selection both strategies are able to take over the population considerably more quickly. Conversely, when the strength of selection is lowered the social strategy is much more effective, as agents are able to build up social capital more effectively when their strategy remains consistent. With a β value of 0.5 the social strategy was able to take over in all 100 simulation runs. Further experiments even showed that with the largest population size of 192 agents, the social strategy was able to take over in 97 of the 100 simulation runs when a β value of 0.5 was used. However, this takeover took a mean of 566 days. Lowering the strength of selection, and hence the rate with which agents change their strategies, is clearly very successful in promoting social behaviour. However, it can greatly slow down the time required for agents to settle on a single strategy.

5 Discussion

We have demonstrated a decentralised and non-monetary mechanism for demand-side management that is effective at satisfying agents' preferences. A real-world implementation of the system could easily operate in a socio-technical manner, with households "in-the-loop" in their agent's decision-making and learning processes [50]. At one level, this could involve households setting the strategy of their agent, based on information presented to them such as their satisfaction metric and current social capital. At a more detailed level, this could involve households interactively making decisions on whether or not to accept requested exchanges. In an actual implementation, the exchanges might occur through software agents that operate on home gateways connected to smart meters, and hosted entirely within households. Alternatively, agents could be hosted on cloud servers where smart meters serve as the intermediary interface [51].

Utilising a system based on social capital represented as ‘favours’ is intuitive for people to understand, agreeing with the reciprocity principle of “I will scratch your back if you scratch mine” [46]. Ease of understanding a mechanism for demand-side management has been shown to be a determinant for whether people perceive the mechanism as procedurally fair, and in turn whether they would be willing to use that mechanism [52]. As such, the use of the reciprocity principle should facilitate perceived procedural justice and hence promote social behaviour within the community. This is in contrast to mechanisms based on dynamic pricing, where evidence from behavioural economics shows that consumers view procedures for setting prices based on supply and demand as unfair [23], which may at least in part explain the low uptake of dynamic pricing tariffs by households.

Future work can investigate the perceived procedural fairness of our mechanism, by inviting participants to use a simulated version of the system [42]. In such a system we can vary the extent to which participants interact with the exchange process, and can control the information that is shown to them, e.g. whether or not they see the satisfaction of other agents, and the number of neutral exchanges that they have performed. The study could also be run through a serious game to help increase participants’ attention and motivation [31]. In these kinds of experiments, the perceived fairness of the mechanism can also be directly compared to other mechanisms such as auctions and dynamic pricing. This approach to investigating perceived procedural fairness should be contrasted with much of the work in the smart energy literature. This literature tends to focus on distributive (outcome) rather than procedural fairness. It typically measures fairness using theoretical measures, such as the mean deviation in agents’ utility from the mean utility [53]. However, measures like this say nothing about the procedure that led to that outcome, suggesting that they are likely to be limited as proxies for household engagement and trust. Consequently, we argue that there is an urgent need to move beyond theoretical measures of fairness and to actually examine perceived fairness in empirical settings.

Beyond the community energy system application, the time-slot exchange mechanism developed here is addressing a classic resource allocation problem – which agents (households) get which resource (time-slots)? Many other approaches to resource allocation, such as the use of virtual auctions, have been found to maximise social welfare – the sum of agents’ utilities – that can be gained from a resource, and can be adapted to also increase the equity of distribution [54, 55]. These approaches are, however, centred around trading some form of currency, which can be either real money or virtual credits [56]. In the case of a decentralised community energy system, even if all households are allocated an equal amount of virtual credits, removing the ability for wealthier individuals to have more bidding power, this approach would still be dependent on competition between agents. Investment and use of a community energy system is reliant on cooperation from a community [10]. We have demonstrated that an approach that is based on incentivising cooperation instead of competition can also be effective. This could allow for stronger social relations to develop if the system is used in a socio-technical manner [37]. Such real-world social capital could, in turn, allow the community to better solve the other cooperation and coordination dilemmas inherent in running a community energy system, such as planning and maintenance [38].

An alternative approach to energy sharing between households is peer-to-peer energy trading [29]. This relies on each house being fitted with its *own* renewable energy source, such as roof-mounted photovoltaic cells. Households with surplus energy can then sell this to households with a deficit of energy, typically using an auction mechanism to determine prices based on supply and demand [57]. Peer-to-peer energy trading thus treats renewable energy sources as individual, private goods, that can be allocated by market mechanisms. A modified version of peer-energy trading can allow households to borrow energy from neighbours and repay that energy when they are able without a direct financial transaction [58]. However, as energy generation will rarely match demand when using renewable energy sources, this approach requires households to be equipped with expensive battery storage as well as their own individual sources of energy generation, making implementation costly for the individual. This cost not only makes communities as a whole less likely to adopt the approach [10], but further risks disadvantaging those who cannot afford the initial cost of household renewable energy generation such as photovoltaic cells and battery storage [59]. Given the importance that access to energy has to both individuals and to the environment, there is a pressing need to consider the concept of *energy justice* when designing mechanisms for the allocation of energy. Energy justice means moving towards an energy system that considers both the

traditional economic needs of societies, but also the environmental issues surrounding climate change and social justice considerations for end users, such as reducing energy poverty [60–62]. Our approach does not require each household to generate energy individually, with the costs that this entails, and instead treats the community's renewable energy sources as a common pool resource [38]. As such, we suggest that this offers more alignment with the principles of energy justice, compared to mechanisms such as peer-to-peer energy trading.

A significant limitation of our simulation model lies in its binary treatment of agent preferences, where they are either fully satisfied or entirely dissatisfied with a given time-slot. In real-world scenarios, users are more likely to retain some level of preference for time-slots near their initial choices. As a result, future research should consider the dynamic impact on the behaviours of both social and selfish agents as the satisfaction they derive from an allocated time-slot becomes contingent on its proximity to their initial preference. This will enable the exploration of more intricate social dynamics, including instances where agents relinquish a moderately satisfying time-slot to benefit others who could attain a higher level of satisfaction. Additionally, this approach will facilitate the examination of scenarios wherein agents may choose to giveaway a time-slot that they somewhat desire (rather than a neutral one) in exchange for the possibility of securing a more favourable time-slot in subsequent exchanges. We also showed that by lowering the strength of selection used when agents are considering changing their strategy, social agents can build up more social capital and the social strategy takes over the population much more consistently. However, in a real implementation the strength of selection would depend on how likely households were to change their agent's strategy. Future work should empirically investigate this, for example, by running user experiments with a simulated version of the system [42]. Future modelling work could also investigate the effect of other agent decision-making mechanisms beyond payoff-biased learning, such as mechanisms based on trust, risk aversion, or social norms.

In conclusion, we have shown how realistic demand peaks make it substantially harder for social agents to take over a population consistently when compared to the artificial flat demand curve used in previous work [13, 14]. However, we have also shown that there is a clear improvement in performance when a population consists of mixed demographics, and so this non-monetary approach to demand-side management can be effective in some realistic scenarios. Our analysis also demonstrates how it is important that the population size is large enough to allow for a high potential mean satisfaction, but not so large that the social agents are unable to build up social capital. One approach for larger populations could be to use a global social capital mechanism – a form of indirect reciprocity [63] – where favours can be repaid by any other agent. However, this could remove the user privacy that comes from agents tracking social capital in a pairwise manner and storing it themselves. As an alternative, larger populations could be grouped into specific community clusters, with a variety of demographics with complementary time-slot preferences. Allowing agents to form self-organised clusters working to optimise their collective performance has been shown to be an effective approach to supply and demand matching within large-scale energy systems [64]. The process of clustering and the subsequent inter-cluster communication will play a vital role in enabling the effectiveness of this demand-side management strategy within real-world markets. This mechanism opens the possibility for distinct clusters to engage in energy trading when surplus energy is available within one cluster [65]. For example, available but unwanted time-slots could be exchanged between clusters, using the reciprocal exchange mechanism analysed in this paper at the level of clusters as well as at the level of households.

Acknowledgements The authors would like to thank the anonymous reviewers for their feedback. The authors used the AI tool ChatGPT (OpenAI) during the revision process, in response to reviewer comments, to review and evaluate the Introduction and Background sections and to suggest improvements to their structure and style. Some AI-suggested structural edits and a small number of AI-generated sentences were incorporated in the Introduction and Background to improve clarity. All content was reviewed and edited by the authors, who take full responsibility for the final version of the manuscript.

Author Contributions J.M.B, N.A.B and S.T.P conceived the study and developed the methodology. N.A.B. wrote the simulation code, analysed the results, and produced all figures and tables. J.M.B. and S.T.P. jointly supervised the research. N.A.B. and S.T.P. wrote the initial draft of the manuscript. J.M.B., N.A.B. and S.T.P. edited the manuscript.

Funding The authors declare that no funding was received from any organisation or agency in support of this research.

Data Availability The raw Household Electricity Survey data analysed during the current study is available from the UK Data Service on request, and a summary report is available at: <https://www.gov.uk/government/publications/household-electricity-survey-2>. The demand curves that we computed from this data are available in the GitHub repository of this paper, which contains full source code for the simulation: <https://github.com/NathanABrooks/EnergyExchangeArena>.

Declarations

Conflict of interest The authors declare no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Skidmore, C.: The Climate Change Act 2008 (2050 Target Amendment) Order 2019. Department for Business, Energy and Industrial Strategy, UK (2019)
2. Sassoli, D.M., Matos Fernandes, J.P.: Regulation (EU) 2021/1119 of the European Parliament and of the Council. European Union (2021)
3. Energy Security, D., Zero, N.: UK energy in brief 2023. National Statistics (2023). https://assets.publishing.service.gov.uk/media/664c827ff34f9b5a56adcb5d/UK_Energy_in_Brief_2023.pdf
4. Walker, G., Devine-Wright, P.: Community renewable energy: What should it mean? *Energy Pol.* **36**(2), 497–500 (2008). <https://doi.org/10.1016/j.enpol.2007.10.019>
5. Gruber, L., Bachhiesl, U., Wogrin, S.: The current state of research on energy communities. *e & i Elektrotechnik und Informationstechnik* **138**(8), 515–524 (2021). <https://doi.org/10.1007/s00502-021-00943-9>
6. Busch, H., Ruggiero, S., Isakovic, A., Hansen, T.: Policy challenges to community energy in the EU : a systematic review of the scientific literature. *Renew. Sustain. Energy Rev.* **151**, 111535 (2021). <https://doi.org/10.1016/j.rser.2021.111535>
7. Chiradeja, P., Ramakumar, R.: An approach to quantify the technical benefits of distributed generation. *IEEE Trans. Energy Convers.* **19**(4), 764–773 (2004). <https://doi.org/10.1109/TEC.2004.827704>
8. Pepermans, G., Driesen, J., Haeseldonckx, D., Belmans, R., D'haeseleer, W.: Distributed generation: definition, benefits and issues. *Energy Pol.* **33**(6), 787–798 (2005). <https://doi.org/10.1016/j.enpol.2003.10.004>
9. Alanne, K., Saari, A.: Distributed energy generation and sustainable development. *Renew. Sustain. Energy Rev.* **10**(6), 539–558 (2006). <https://doi.org/10.1016/j.rser.2004.11.004>
10. Pasimeni, F.: Community-based adoption and diffusion of micro-grids: analysis of the Italian case with agent-based model. *Journal of Artificial Societies and Social Simulation* (2019). <https://doi.org/10.18564/jasss.3950>
11. Kanakadhurga, D., Prabakaran, N.: Demand side management in microgrid: a critical review of key issues and recent trends. *Renew. Sustain. Energy Rev.* **156**, 111915 (2022). <https://doi.org/10.1016/j.rser.2021.111915>
12. Warren, P.: A review of demand-side management policy in the UK. *Renew. Sustain. Energy Rev.* **29**, 941–951 (2014). <https://doi.org/10.1016/j.rser.2013.09.009>
13. Petruzzzi, P.E., Busquets, D., Pitt, J.: Self organising flexible demand for smart grid. In: 2013 IEEE 7th International Conference on Self-Adaptation and Self-Organizing Systems Workshops, IEEE, pp. 21–22 (2013). <https://doi.org/10.1109/SASOW.2013.19>
14. Brooks, N.A., Powers, S.T., Borg, J.M.: A mechanism to promote social behaviour in household load balancing. In: ALIFE 2020: The 2020 Conference on Artificial Life, pp. 95–103 (2020). https://doi.org/10.1162/isal_a_00290
15. Zimmermann, J.-P., Evans, M., Griggs, J., King, N., Harding, L., Roberts, P., Evans, C.: Household electricity survey: a study of domestic electrical product usage. Intertek Testing & Certification Ltd, 213–214 (2012)
16. Stern, P.C., Aronson, E., Darley, J.M., Hill, D.H., Hirst, E., Kempton, W., Wilbanks, T.J.: The effectiveness of incentives for residential energy conservation. *Eval. Rev.* **10**(2), 147–176 (1986). <https://doi.org/10.1177/0193841X8601000201>
17. Dutta, G., Mitra, K.: A literature review on dynamic pricing of electricity. *J. Oper. Res. Soc.* **68**(10), 1131–1145 (2017). <https://doi.org/10.1057/s41274-016-0149-4>

18. Yuniarto, H.A., Normasari, N.M.E., Silalahi, S.F., Gunawan, I.C., Aditya, I.A., Simaremare, A.A., Haryadi, F.N.: Customer behaviour towards energy usage with time of use tariff: a systematic literature review. *J. Revenue Pric. Manag.* **23**(1), 44–61 (2024)
19. Hao, C.H., Wesseh, P.K., Wang, J., Abudu, H., Dogah, K.E., Okorie, D.I., Osei Opoku, E.E.: Dynamic pricing in consumer-centric electricity markets: a systematic review and thematic analysis. *Energ. Strat. Rev.* **52**, 101349 (2024). <https://doi.org/10.1016/j.esr.2024.101349>
20. Panda, S., Mohanty, S., Rout, P.K., Sahu, B.K., Parida, S.M., Samanta, I.S., Bajaj, M., Piecha, M., Blazek, V., Prokop, L.: A comprehensive review on demand side management and market design for renewable energy support and integration. *Energy Rep.* **10**, 2228–2250 (2023). <https://doi.org/10.1016/j.egyr.2023.09.049>
21. Junaidi, N., Abdullah, M.P., Alharbi, B., Shaaban, M.: Blockchain-based management of demand response in electric energy grids: a systematic review. *Energy Rep.* **9**, 5075–5100 (2023). <https://doi.org/10.1016/j.egyr.2023.04.020>
22. Azarova, V., Engel, D., Ferner, C., Kollmann, A., Reichl, J.: Exploring the impact of network tariffs on household electricity expenditures using load profiles and socio-economic characteristics. *Nat. Energy* **3**(4), 317–325 (2018). <https://doi.org/10.1038/s41560-018-0105-4>
23. Kahneman, D., Knetsch, J.L., Thaler, R.: Fairness as a constraint on profit seeking: entitlements in the market. *Am. Econ. Rev.* **76**(4), 728–741 (1986)
24. Neuteleers, S., Mulder, M., Hindriks, F.: Assessing fairness of dynamic grid tariffs. *Energy Pol.* **108**, 111–120 (2017). <https://doi.org/10.1016/j.enpol.2017.05.028>
25. Gneezy, U., Rustichini, A.: A fine is a price. *J. Leg. Stud.* **29**(1), 1–17 (2000). <https://doi.org/10.1086/468061>
26. Promberger, M., Marteau, T.M.: When do financial incentives reduce intrinsic motivation? Comparing behaviors studied in psychological and economic literatures. *Health Psychol.* **32**(9), 950 (2013). <https://doi.org/10.1037/a0032727>
27. Burson, J., Harvey, N.: Mo money, Mo problems: when and why financial incentives backfire. *J. Manag. Inf. Decis. Sci.* **22**(3), 191–206 (2019)
28. Ostrom, E.: *Governing the Commons : The Evolution of Institutions for Collective Action*. Cambridge University Press, Cambridge, UK (1990)
29. Tushar, W., Yuen, C., Saha, T.K., Morstyn, T., Chapman, A.C., Alam, M.J.E., Hanif, S., Poor, H.V.: Peer-to-peer energy systems for connected communities: a review of recent advances and emerging challenges. *Appl. Energy* **282**, 116131 (2021). <https://doi.org/10.1016/j.apenergy.2020.116131>
30. DellaValle, N., Sareen, S.: Nudging and boosting for equity? Towards a behavioural economics of energy justice. *Energy Res. Soc. Sci.* **68**, 101589 (2020). <https://doi.org/10.1016/j.erss.2020.101589>
31. Bourazeri, A., Pitt, J., Arnab, S.: Enabling collective awareness of energy use via a social serious game. *EAI Endorsed Trans. Game Based Learn.* **4**(13), 153510 (2017). <https://doi.org/10.4108/eai.27-12-2017.153510>
32. Khanna, T.M., Baiocchi, G., Callaghan, M., Creutzig, F., Guias, H., Haddaway, N.R., Hirth, L., Javaid, A., Koch, N., Laukemper, S., Zamora, M.D.M., Minx, J.C.: A multi-country meta-analysis on the role of behavioural change in reducing energy consumption and CO2 emissions in residential buildings. *Nat. Energy* **6**(9), 925–932 (2021). <https://doi.org/10.1038/s41560-021-00866-x>
33. Vesely, S., Klöckner, C.A., Carrus, G., Tiberio, L., Caffaro, F., Biresselioglu, M.E., Kollmann, A.C., Sinea, A.C.: Norms, prices, and commitment: a comprehensive overview of field experiments in the energy domain and treatment effect moderators. *Front. Psychol.* (2022). <https://doi.org/10.3389/fpsyg.2022.967318>
34. Fadhuile, A., Llerena, D., Roussillon, B.: Intrinsic motivation to promote demand flexibility: a field experiment from household demand. *Energy J.* (2025). <https://doi.org/10.1177/01956574251320363>
35. Jorgensen, B.S., Fumei, S., Byrne, G.: Reducing peak energy demand among residents who are not billed for their electricity consumption: experimental evaluation of behaviour change interventions in a university setting. *Int. J. Environ. Res. Public Health* **18**(16), 8406 (2021). <https://doi.org/10.3390/ijerph18168406>. Number: 16 Publisher: Multidisciplinary Digital Publishing Institute
36. Streimikiene, D.: Use of nudges for promotion of sustainable energy consumption in households. *Contemp. Econ.* **17**(1), 1–9 (2023). <https://doi.org/10.5709/ce.1897-9254.495>
37. Bellman, K., Botev, J., Hildmann, H., Lewis, P.R., Marsh, S., Pitt, J., Scholtes, I., Tomforde, S.: Socially-sensitive systems design: exploring social potential. *IEEE Technol. Soc. Mag.* **36**(3), 72–80 (2017). <https://doi.org/10.1109/MTS.2017.2728727>
38. Pitt, J.: *Self-Organising Multi-Agent Systems: Algorithmic Foundations of Cyber-Anarcho-Socialism*. World Scientific (Europe), London (2021). <https://doi.org/10.1142/q0307>
39. Fina, B., Auer, H., Friedl, W.: Profitability of contracting business cases for shared photovoltaic generation and renovation measures in a residential multi-apartment building. *J. Clean. Prod.* **265**, 121549 (2020). <https://doi.org/10.1016/j.jclepro.2020.121549>
40. Pasina, A., Canoilas, A., Johansson, D., Bagge, H., Fransson, V., Davidsson, H.: Shared PV systems in multi-scaled communities. *Buildings* **12**(11), 1846 (2022). <https://doi.org/10.3390/buildings12111846>
41. Galvin, R.: How photovoltaics make energy refurbishment more affordable in apartment buildings. *J. Clim. Finance* **7**, 100039 (2024). <https://doi.org/10.1016/j.jclimf.2024.100039>

42. Powers, S.T., Meanwell, O., Cai, Z.: Finding fair negotiation algorithms to reduce peak electricity consumption in micro grids. In: International Conference on Practical Applications of Agents and Multi-Agent Systems, pp. 269–272 (2019). https://doi.org/10.1007/978-3-030-24209-1_28. Springer
43. Putnam, R.D.: Social capital and public affairs. *Bulletin of the American Academy of Arts and Sciences*, 5–19 (1994) <https://doi.org/10.2307/3824796>
44. Petruzzi, P.E., Busquets, D., Pitt, J.: Experiments with social capital in multi-agent systems. In: International Conference on Principles and Practice of Multi-Agent Systems, Springer, pp. 18–33 (2014). https://doi.org/10.1007/978-3-319-13191-7_2
45. Kua, J., Hossain, M.B., Natgunanathan, I., Xiang, Y.: Privacy preservation in smart meters: current status. *Challeng. Fut. Direct. Sens.* **23**(7), 3697 (2023). <https://doi.org/10.3390/s23073697>
46. Hoffman, E., McCabe, K.A., Smith, V.L.: Behavioral foundations of reciprocity: experimental economics and evolutionary psychology. *Econ. Inq.* **36**(3), 335–352 (1998). <https://doi.org/10.1111/j.1465-7295.1998.tb01719.x>
47. Powers, S.T., Lehmann, L.: When is bigger better? The effects of group size on the evolution of helping behaviours. *Biol. Rev.* **92**, 902–920 (2017). <https://doi.org/10.1111/brv.12260>
48. Traulsen, A., Pacheco, J.M., Nowak, M.A.: Pairwise comparison and selection temperature in evolutionary game dynamics. *J. Theor. Biol.* **246**(3), 522–529 (2007). <https://doi.org/10.1016/j.jtbi.2007.01.002>
49. Palmer, J., Terry, N., Kane, T.: Further analysis of the household electricity survey - Early findings: Demand side management. Loughborough University and Element Energy, London, United Kingdom, Cambridge Architectural Research Limited (2013)
50. Pitt, J., Diaconescu, A., Bourazeri, A.: Democratisation of the SmartGrid and the active participation of prosumers. In: 2017 IEEE 26th International Symposium on Industrial Electronics (ISIE), pp. 1707–1714 (2017). <https://doi.org/10.1109/ISIE.2017.8001505>. ISSN: 2163-5145
51. Saoud, A., Recioui, A.: Hybrid algorithm for cloud-fog system based load balancing in smart grids. *Bull. Electr. Eng. Inf.* **11**(1), 477–487 (2022)
52. Schlereth, C., Skiera, B., Schulz, F.: Why do consumers prefer static instead of dynamic pricing plans? An empirical study for a better understanding of the low preferences for time-variant pricing plans. *Eur. J. Oper. Res.* **269**(3), 1165–1179 (2018). <https://doi.org/10.1016/j.ejor.2018.03.033>
53. Tsousoglou, G., Efthymiopoulos, N., Makris, P., Varv Arigos, E.: Personalized real time pricing for efficient and fair demand response in energy cooperatives and highly competitive flexibility markets. *J. Mod. Power Syst. Clean Energy* **7**(1), 151–162 (2019). <https://doi.org/10.1007/s40565-018-0426-0>
54. Buyya, R., Abramson, D., Giddy, J., Stockinger, H.: Economic models for resource management and scheduling in grid computing. *Concurr. Comput. Pract. Exp.* **14**(13–15), 1507–1542 (2002). <https://doi.org/10.1002/cpe.690>
55. Pla, A., Lopez, B., Murillo, J.: Multi-dimensional fairness for auction-based resource allocation. *Knowl.-Based Syst.* **73**, 134–148 (2015). <https://doi.org/10.1016/j.knosys.2014.09.009>
56. Zhang, Y., Lee, C., Niyato, D., Wang, P.: Auction approaches for resource allocation in wireless systems: a survey. *IEEE Commun. Surv. Tutor.* **15**(3), 1020–1041 (2012). <https://doi.org/10.1109/SURV.2012.110112.00125>
57. Teixeira, D., Gomes, L., Vale, Z.: Single-unit and multi-unit auction framework for peer-to-peer transactions. *Int. J. Electr. Power Energy Syst.* **133**, 107235 (2021). <https://doi.org/10.1016/j.ijepes.2021.107235>
58. Prasad, A., Dusparic, I.: Multi-agent deep reinforcement learning for zero energy communities. In: 2019 IEEE PES Innovative Smart Grid Technologies Europe (ISGT-Europe), IEEE, pp. 1–5 (2019). <https://doi.org/10.1109/ISGTEurope.2019.8905628>
59. Kim, J.D.-Y., Choi, E., Lee, E.: Social justice, fairness and exclusion in the South Korean electricity sector. *Energy Res. Soc. Sci.* **51**, 55–66 (2019). <https://doi.org/10.1016/j.erss.2018.12.002>
60. Heffron, R.J., McCauley, D.: The concept of energy justice across the disciplines. *Energy Pol.* **105**, 658–667 (2017). <https://doi.org/10.1016/j.enpol.2017.03.018>
61. Jenkins, K., McCauley, D., Heffron, R., Stephan, H., Rehner, R.: Energy justice: a conceptual review. *Energy Res. Soc. Sci.* **11**, 174–182 (2016). <https://doi.org/10.1016/j.erss.2015.10.004>
62. Sovacool, B.K., Dworkin, M.H.: Energy justice: conceptual insights and practical applications. *Appl. Energy* **142**, 435–444 (2015). <https://doi.org/10.1016/j.apenergy.2015.01.002>
63. Sigmund, K.: *The Calculus of Selfishness*. Princeton University Press, Princeton (2010)
64. Čaušević, S., Warnier, M., Brazier, F.M.: Dynamic, self-organized clusters as a means to supply and demand matching in large-scale energy systems. In: 2017 IEEE 14th International Conference on Networking, Sensing and Control (ICNSC), pp. 568–573 (2017). <https://doi.org/10.1109/ICNSC.2017.8000154>
65. Cardoso, R.P., Hart, E., Pitt, J.V.: Evolving robust policies for community energy system management. In: Proceedings of the Genetic and Evolutionary Computation Conference. GECCO '19, pp. 1120–1128. ACM, New York, NY, USA (2019). <https://doi.org/10.1145/3321707.3321763>

Authors and Affiliations

Nathan A. Brooks¹ · Simon T. Powers² · James M. Borg³

✉ Simon T. Powers
s.t.powers@stir.ac.uk

Nathan A. Brooks
n.a.brooks@keele.ac.uk

James M. Borg
j.borg@aston.ac.uk

¹ School of Computer Science and Mathematics, Keele University, Staffordshire, UK

² Division of Computing Science and Mathematics, University of Stirling, Stirling, UK

³ School of Computer Science and Digital Technologies, Aston University, Birmingham, UK