

A mechanism to promote social behaviour in household load balancing

Nathan A. Brooks^{1,*}, Simon T. Powers^{2,3} and James M. Borg¹

¹ School of Computing and Mathematics, Keele University, Staffordshire, UK, ST5 5BG

² School of Computing, Edinburgh Napier University, Edinburgh, UK, EH10 5DT

³ Institute of Liberal Arts and Sciences, Keele University, Staffordshire, UK, ST5 5BG

* n.a.brooks@keele.ac.uk

Abstract

Reducing the peak energy consumption of households is essential for the effective use of renewable energy sources, in order to ensure that as much household demand as possible can be met by renewable sources. This entails spreading out the use of high-powered appliances such as dishwashers and washing machines throughout the day. Traditional approaches to this problem have relied on differential pricing set by a centralised utility company. But this mechanism has not been effective in promoting widespread shifting of appliance usage. Here we consider an alternative decentralised mechanism, where agents receive an initial allocation of time-slots to use their appliances and can then exchange these with other agents. If agents are willing to be more flexible in the exchanges they accept, then overall satisfaction, in terms of the percentage of agents' time-slot preferences that are satisfied, will increase. This requires a mechanism that can incentivise agents to be more flexible. Building on previous work, we show that a mechanism incorporating social capital — the tracking of favours given and received — can incentivise agents to act flexibly and give favours by accepting exchanges that do not immediately benefit them. We demonstrate that a mechanism that tracks favours increases the overall satisfaction of agents, and crucially allows social agents that give favours to outcompete selfish agents that do not under payoff-biased social learning. Thus, even completely self-interested agents are expected to learn to produce socially beneficial outcomes.

Introduction

The UK government has committed to a legally binding target to reduce greenhouse gas emissions by 80% by 2050. This requires a shift to renewable energy sources, such as solar panels and wind turbines. However, integrating renewable energy sources into a centralised and monolithic 'national grid' is difficult because their output inherently depends on weather conditions. As such, they cannot simply be 'switched on and off' to meet demand in the way that coal, gas, and nuclear power stations can be to match supply and demand. Governments and energy providers have recognised that this problem of load balancing – matching supply and demand – is easier to solve on a local scale. Consequently, they are supporting the development of 'community energy systems', where a community (e.g. a town or a

small island) owns and manages its own renewable energy sources (Walker and Devine-Wright, 2008).

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. A key problem here 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.

The traditional approach to reducing peak demand is differential pricing set by a central utility company. Simply put, households are incentivised to run their appliances at times of low demand through lower pricing at these times (Stern et al., 1986; Dutta and Mitra, 2017). Traditionally, this has involved utility companies offering cheaper electricity overnight. Could a community energy system use the same mechanism? Potentially, however, variable pricing inherently discriminates against more vulnerable households on lower incomes (Simmons and Rowlands, 2007). Then there is the question of how prices should be set and who should set them? People are unlikely to take part in such a scheme unless they perceive that they are being treated fairly.

To address these issues, we consider an alternative mechanism for load balancing in a community energy system, which is not based on pricing set by some centralised authority. We assume that each household will have preferred time-slots for when they would like to run high-powered appliances such as washing machines, dishwashers and electric heating. The aim is then to allocate actual time-slots to each household for when they run their appliances. On the one hand, this is a classic multi-objective optimisation problem of reducing peak consumption (the maximum amount of energy demanded in a time-slot) while satisfying each household's preferences as far as possible. On the other hand, issues of fairness are central. If households are to be moti-

vated to use the mechanism then they will need to perceive the resulting *allocation* of time-slots to households as being fair (distributive justice, Rescher 1966). Furthermore, households will need to be able to understand why some of their time-slot preferences have not been met, and why the preferences of some other households may have been met instead. In other words, they need to perceive the allocation *procedure* as treating them fairly (procedural justice, Hollander-Blumoff and Tyler 2008).

Petruzzi et al. (2013) propose a mechanism inspired by the building of social capital between agents (households, or software agents representing them). In their mechanism, agents are initially allocated time-slots at random, but can then propose exchanges of time-slots with other agents. Agents have two possible strategies. Selfish agents only accept exchanges that provide them with one of their preferred time-slots. Social agents, on the other hand, accept not only these beneficial exchanges but also accept an exchange request if they owe a favour to another agent (provided the exchange will not cause them to lose one of their preferred time-slots). An agent owes a favour to another if the other agent has previously accepted an exchange request from them. Petruzzi et al. (2013) showed that under this mechanism, a group where every agent was social had on average more of their time-slot preferences satisfied than a group where every agent was selfish. They construed the recording of favours given and received, and the acting upon this by social agents, as the accumulation of a form of electronic social capital (Putnam, 1994; Petruzzi et al., 2014). Presumably, this would be intuitive for households to understand.

However, two important questions arise from this work. First, to what extent is social capital a necessary part of the mechanism? Would social agents that always accept exchanges that do not cause them to lose a preferred time-slot also reach outcomes with high average satisfaction, without the need to track favours given and received? Second, should we expect self-interested agents to adopt the selfish or the social strategy? We address both of these questions in this paper. To address the first, we re-examine the Petruzzi et al. (2013) model by allowing social capital to be turned off. To address the second, we consider mixed populations of selfish and social agents that change their strategy according to payoff-biased social learning (Boyd and Richerson, 1985). We find that while a mechanism without social capital allows a pure population of social agents to perform better than a pure population of selfish agents, when agents can change their strategy through social learning then social capital is necessary for social agents to outcompete selfish agents. Our results demonstrate that a time-slot allocation mechanism based on social capital can reduce peak electricity consumption, promote social behaviour, and lead to outcomes where the average satisfaction of households is high, even when agents are entirely self-interested.

The Energy Exchange Simulation

The Petruzzi et al. (2013) model was built to represent a smart energy network consisting of 96 individual agents. Each day agents request four hour-long time-slots in which they require electricity. All requests are for 1KWh of energy, and there can never be more than 16 agents using the same time-slot, as this is considered the peak capacity of the system. Time-slots are initially allocated at random at the start of the day, so few agents are likely to have their allocation match all of their requested time-slots. Because of this, after the initial allocation agents can partake 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. We define an agent's *satisfaction* as the proportion of its time-slot preferences that have been satisfied, and track the mean value of this as a measure of how well the mechanism is satisfying the agents' preferences.

Agents can follow either social or selfish strategies, which impacts how they react to incoming requests for exchanges. 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 and do not already have in order for them to agree to the exchange. Social agents also agree to these mutually beneficial exchanges. However, they 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 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 is not detrimental to their own satisfaction and record that the favour has been repaid. This improves the satisfaction of the other agent while earning themselves more social capital. This leads to a system of social agents earning and repaying favours among one another, increasing the number of accepted exchange requests.

Exchanges begin on a 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'. Several exchange rounds then take place during the day, where the number of rounds is a parameter of the model that sets the maximum number of exchanges an agent can engage in per day. 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. Agents accept or refuse requests based on their strategy as described above. Only social capital, i.e. social agents' memory of favours, remains between days.

We expand on the original Petruzzi et al. (2013) model by introducing social learning, allowing agents to change from selfish to social or vice versa (note that both social and selfish agents undergo 'social' learning, which we re-

fer to simply as ‘learning’ from now on to avoid confusion). This works as follows. At the end of each day, a percentage of the agents are randomly selected to undergo learning. Each agent performing learning observes a randomly selected second agent. If the observed agent has a higher satisfaction than the agent in question, then the first agent will copy the strategy of the observed agent with a probability proportionate to the difference between the two agents’ satisfactions. Learning is thus payoff-biased (Boyd and Richerson, 1985), with strategies giving higher individual satisfaction more likely to spread in the population. Agents that move from a social strategy to the selfish strategy retain their accumulated social capital. Pseudocode for the simulation procedure is given in Algorithm 1¹.

We also consider, as a counterpoint to social capital, simulations where social capital is not recorded by agents, resulting in social agents accepting any non-detrimental exchange. We refer to the social agents under this mechanism as *social without social capital*. All simulation results are averaged over 50 runs for each variation of the simulation parameters. The parameters that we vary are the number of exchange rounds per day, and the number of agents undergoing social learning at the end of each day (the learning rate). We hold the other parameters constant across simulations with the values given in Table 1.

Parameter	Value
Population size	96
Number of days	500
Time-slots per day	24
Slots selected by each agent	4
Maximum agents per time-slot	16
Simulation runs	50

Table 1: Constant parameter values.

Results from an illustrative run of the Energy Exchange Simulation can be seen in Figure 1, showing how agent satisfaction changes over time.

Results and Analysis

To answer the research questions raised in the introduction, our analysis of the simulation proceeds as follows. We first consider populations where every agent uses the same strategy, and there is no learning. This allows us to determine whether a pure population of social agents using the mechanism without social capital can do as well as a pure population of agents that do track social capital. We then go on to consider mixed populations where both selfish and social agents are present and can switch their strategies through learning. This allows us to determine the conditions under which social agents can outcompete selfish agents, and the

¹Source code is available at <https://github.com/NathanABrooks/ResourceExchangeArena>

Algorithm 1 The Energy Exchange Simulation.

```

1:  $d \leftarrow$  current day
2:  $e \leftarrow$  current exchange_round
3:  $A \leftarrow$  set of  $a$  agents
4:  $L \leftarrow$  number of agents  $a$  undergoing learning
5: for  $d = 1$  to MAX_DAYS do
6:   for each  $a \in A$  do
7:      $a.receive\_random\_allocation()$ 
8:   end for
9:   for  $e = 1$  to MAX_EXCHANGES do
10:     $V \leftarrow$  set of  $v$  adverts
11:    for each  $a \in A$  do
12:       $v \leftarrow a.determine\_unwanted\_time\_slots()$ 
13:       $V.list\_advert(v)$ 
14:    end for
15:    for each  $a \in A$  do
16:      if  $a.received\_request() == \text{true}$  then
17:        go to next agent
18:      end if
19:      if  $a.satisfaction() == 1$  then
20:        go to next agent
21:      else
22:         $r \leftarrow a.identify\_beneficial\_exchange(V)$ 
23:         $a.request\_exchange(r)$ 
24:      end if
25:    end for
26:    for each  $a \in A$  do
27:      if  $a.received\_request() == \text{true}$  then
28:         $a.accept\_exchange\_if\_approved()$ 
29:      end if
30:    end for
31:    for each  $a \in A$  do
32:      if  $a.made\_exchange()$  and  $a.agent\_type == \text{So-}$ 
33:         $a.update\_social\_capital()$ 
34:      end if
35:    end for
36:  end for
37:   $d_l \leftarrow 0$ 
38:  while  $d_l < L$  do
39:     $a_1 \leftarrow$  random agent that hasn’t considered chang-
40:    ing strategy today
41:     $a_2 \leftarrow$  random agent to observe
42:    if  $a_1.satisfaction() < a_2.satisfaction()$  then
43:       $x \leftarrow$  random value between 0 and 1
44:      if  $(a_2.satisfaction() - a_1.satisfaction()) > x$  then
45:         $a_1.copy\_strategy(a_2)$ 
46:      end if
47:    end if
48:     $d_l \leftarrow d_l + 1$ 
49:  end while
50: end for

```




Figure 1: Average satisfaction per day for a typical run of the simulation. The run shown here includes social capital, uses a learning rate of 50%, and 100 exchange rounds per day. The optimum result is calculated as the proportion of requested time-slots that exist within the simulation and could therefore be allocated to an agent who requested them.

extent to which this is affected by whether social capital is tracked.

Single Strategy Populations

To establish a baseline for the performance of each of the available strategies explored here (selfish, and social with and without social capital), each is first run in isolation of the others, i.e. as a pure population of the strategy. Average (mean) satisfaction heatmaps for these simulations, across all parameters, can be seen in Figure 2. Optimal performance in these simulations would result in an average satisfaction of approximately 0.91 (as seen in Figure 1). From Figure 2 we can see that social populations with social capital achieve this optimal performance over the vast majority of parameters settings. Selfish populations, on the other hand, consistently fall short of optimal satisfaction, but do show improvement as the number of rounds of exchanges on a day increases. No improvement is seen as the number of days the simulation is run for is increased, which follows from the fact that the only state carried over between days is social capital, with selfish agents do not use. Social populations without the social capital mechanism similarly exhibit no improvement as the number of days increases for the same reason. But crucially, they reach optimal performance on the first day, if afforded enough exchanges. The inclusion of social capital means that agents require more days to reach (near) optimal performance, with this only being achieved after 100 days of social capital accumulation.

Overall, it is clear from Figure 2 that satisfaction is dependent on the number of exchanges available to agents, and that social agents are able to break out sub-optimal

states more easily than selfish agents. This is through social agents accepting exchanges that are neutral to their satisfaction but that increase the satisfaction of another agent. However, the tracking and consideration of social capital slows this process down compared to the mechanism without social capital. This is because with tracking of social capital, social agents will only accept neutral exchanges if they owe a favour to the other agent, and so they behave like selfish agents until they begin to owe favours from previous days.

Mixed Populations without Social Capital

Results for populations which combine selfish and social agents, but without access to social capital (resulting in social agents always accepting non-detrimental exchanges) can be seen in Figure 3. In these simulations, populations start with an equal number of social and selfish agents, with learning taking place at the end of each day, affecting the overall ratio of social to selfish agents. Different learning rates are also tested; 0%, resulting in all agents retaining the strategy they were initialised with, 50%, permitting half of the agents (at random) to undergo learning, and 100%, resulting in all agents engaging in learning per day. Figure 3 (Bottom Row) shows the change in strategy distribution across all parameters and learning rates. Within the maximum number of days available in these simulations neither the social nor selfish strategy is eradicated from the population (which is possible with the payoff-biased learning in Algorithm 1). In fact, it is often the case that close to a 50:50 distribution of strategies is retained, indicating no significant change due to learning. We do find that the social strategy begins to dominate the population when 50 exchanges per day are allowed, the simulation has run for over half the maximum number of days, and learning is implemented. But when the learning rate is set to 50%, and the number of exchanges exceeds 150, we see the population distributions beginning to swing in the favour of the selfish strategy. Increasing the learning rate to 100% does stop the selfish strategy gaining prominence under these conditions but does not even get close to eradicating the selfish strategy.

Given the (near) optimal performance demonstrated by social agents (without social capital) observed in Figure 2, where the social strategy existed in isolation, it is clear from Figure 3 (Middle Row) that the inclusion of the selfish strategy has a detrimental effect of the satisfaction of social agents, whilst improving the satisfaction of selfish agents (Figure 3 (Top Row)). The only time the social strategy achieves an advantage over the selfish strategy is when exchanges are set at 50 per day – this result appears to be more a result of 50 exchanges not being sufficient to allow the selfish strategy to achieve its best performance, whilst being sufficient for the social strategy to do so. The negative consequences of the continued use of the selfish strategy to overall satisfaction of the populations, compared to when so-

Average Satisfaction in Single Strategy Populations




Figure 2: Average satisfaction of single strategy populations. For all plots the y -axis shows the simulated day satisfaction was measured on, and the x -axis shows the number of exchanges an agent is permitted to engage in per day. (Left) Average satisfaction for populations of selfish agents. (Middle) Average satisfaction for populations of social agents without social capital. (Right) Average satisfaction for populations of social agents with social capital.

cial strategies are applied in isolation, are clearly apparent. Whilst the social strategy is necessary for selfish agents to improve their satisfaction, through exploitation of the social agents’ willingness to exchange, it is to the detriment of the overall satisfaction of the population. It is therefore in the interests of the populations to remove the selfish strategy. In effect, without social capital selfish agents are able to parasitise social agents to the detriment of the population as a whole.

Mixed Populations with Social Capital

As in Figure 2, populations using the social strategy whilst keeping track of social capital exhibit a slower progression toward optimal satisfaction, as they only start accepting neutral exchanges as social capital accumulates over the days. This is also the case when the selfish strategy is included in a combined population (Figure 4). However, unlike mixed populations where social capital is not tracked (see Figure 3), populations tracking and using social capital are able to achieve optimal (or near optimal) satisfaction when afforded enough exchanges, and the simulation is ran for over 100 days. As the learning rate is increased to 50% the number of exchanges required to achieve near optimal average satisfaction scores amongst social agents drops, though it is interesting to note that increasing the learning rate again to 100% does cause a drop in satisfaction when fewer exchanges per day are permitted.

Social capital thus has the effect of slowing down optimisation, but with the benefit of hindering the ability of selfish agents to gain any traction in the population. We can see that when learning is set to 0% selfish agents, despite making up half of the population, struggle achieve the same kinds of satisfaction scores as previously seen when social capital is not available (compare Figure 3 (Left) to Figure 4 (Left)). As these selfish agents cannot accumulate social capital, de-

sired time-slots held by other agents are harder to obtain than when social capital is not tracked. This enables social agents to build up exchange networks using social capital that cannot be invaded by selfish agents. When learning is introduced, the inclusion of social capital effectively eradicates the selfish strategy from the population. As social agents are on average more satisfied than selfish ones, with this effect becoming more apparent as the simulation goes on for more days or more exchanges are permitted, selfish agents increasingly switch strategies. Conversely, with no opportunity for selfish agents to parasitise the population there is little gradient for social agents to switch to the selfish strategy. We can see in Figure 4 (Bottom) that inclusion of learning and social capital results in the entire population adopting the social strategy. Increasing learning to 100% does slow the removal of the selfish strategy down, as it introduces the possibility of social agents switching to the selfish strategy in cases where some selfish agents have achieved high satisfaction by luck alone (i.e. being randomly allocated all of their preferred time slots), but generally learning combined with social capital has the effects of removing the selfish strategy from the population, leading to near optimal satisfaction scores when provided with enough exchanges and enough simulated days.

The Population-level Effect of Social Capital

Despite the introduction of social capital removing the selfish strategy, and achieving near optimal satisfaction across a number of parameter settings, it is still the case the overall population satisfactions (averaged over both selfish and social agents) are similar when comparing populations with and without social capital (Figure 5). Due to social capital slowing down the rate at which social agents undergo exchanges, when only a low number of exchanges per day are available, or the simulations are stopped after a few

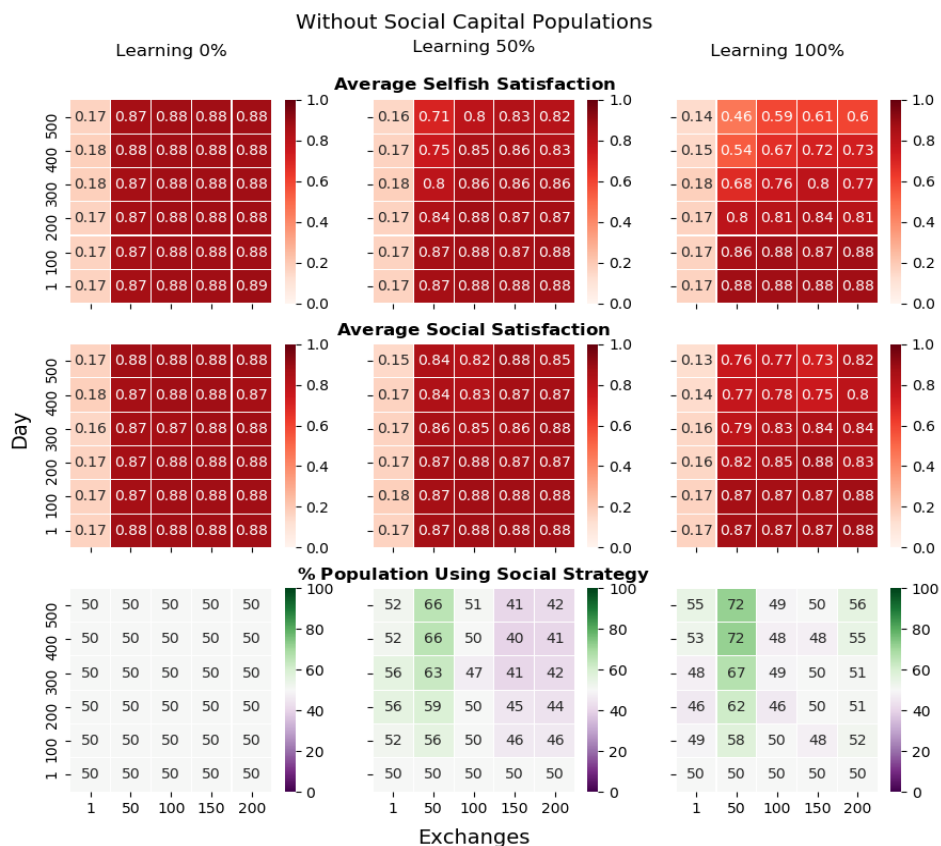


Figure 3: Average satisfaction of selfish and social agents in mixed populations without social capital. For all plots the y -axis shows the simulated day satisfaction was measured on, and the x -axis shows the number of exchanges an agent is permitted to engage in per day. (Top Row) Average satisfaction of selfish agents. (Middle Row) Average satisfaction of social agents. (Bottom Row) Proportion of population using the social strategy; green indicates a greater proportion of social agents; purple indicates a greater proportion of selfish agents.

days, populations without social capital actually outperform populations with social capital, despite selfish agents being largely present in populations without social capital. Once learning is introduced, the number of exchanges per day are increased, and the simulation is permitted to run for extended amount of time, we do begin to see social capital populations outperforming those without social capital, though only by a small margin. Taking just the raw satisfaction results for the 500th day, we do see a significant difference between the satisfaction of social capital and non social capital populations over most parameter settings. Using a Mann-Whitney U test, we observe p -values where $p < 0.01$ over most parameter settings (see Table 2), the exceptions being when learning is 50% and exchanges are set to 1, and when learning is 100% and exchanges are low (1 or 50). These results indicate that by the 500th day of the simulation, social capital populations where learning is enabled (at 50% or above) are significantly more satisfied than non social capital populations.

Mann-Whitney U: $p < 0.01$			
Learning \ Exchanges	0%	50%	100%
1	True	False	False
50	True	True	False
100	True	True	True
150	True	True	True
200	True	True	True

Table 2: Mann-Whitney U test p -values when comparing satisfaction of agents in populations with and without social capital. Test conducted for 500th day only. $p < 0.01 = True$, $p \geq 0.01 = False$.

Discussion

We have extended the [Petruzzi et al. \(2013\)](#) model to answer two research questions. Our first question was, to what extent is the tracking of social capital a necessary feature




Figure 4: Average satisfaction of selfish and social agents in mixed populations with social capital. For all plots the y -axis shows the simulated day satisfaction was measured on, and the x -axis shows the number of exchanges an agent is permitted to engage in per day. (Top Row) Average satisfaction of selfish agents. (Middle Row) Average satisfaction of social agents. (Bottom Row) Proportion of population using the social strategy; green indicates a greater proportion of social agents; purple indicates a greater proportion of selfish agents.




Figure 5: Average difference in satisfaction between all agents in populations with and without social capital. In all rows purple indicates higher satisfaction in populations without social capital, orange indicates higher satisfaction in populations with social capital. For all plots the y -axis shows the simulated day satisfaction was measured on, and the x -axis shows the number of exchanges an agent is permitted to engage in per day.

of the exchange mechanism? We demonstrated that if social agents exist in isolation in a pure population, then they end up more satisfied when social capital is *not* tracked, for cases where the number of days that the simulation is run for is low. When the number of days is increased, they perform similar to when social capital is tracked. Thus, in pure populations with no possibility of selfish agents being introduced then tracking of social capital is not beneficial. Moreover, tracking social capital is detrimental under some conditions since it requires more days before social agents start accepting neutral exchanges, and hence more days before they escape from sub-optimal initial allocations.

However, in a real system we must allow for the possibility that agents may act selfishly, or may learn to do so over time. To account for this, we have considered mixed populations where both social and selfish agents are present and can exchange with each other. In this case, social agents are able to achieve (near to) optimal satisfaction when social capital is tracked, whereas without social capital they do less well. This effect becomes more pronounced once agents are permitted to adjust their strategy by payoff-biased learning. This is because without social capital, selfish agents are able to persist in the population under payoff-biased learning. Their persistence is detrimental both to social agents and to the population as a whole, since by not accepting neutral exchanges they can prevent the agents escaping from sub-optimal allocations. Conversely, with social capital selfish agents are effectively purged from the population under payoff-biased learning, which has a significant positive effect on the satisfaction of the collective population.

The results with payoff-biased (social) learning also answer our second research question: should we expect self-interested agents to adopt the selfish or the social strategy? This is because under payoff-biased learning, agents will only change their strategy if they see another agent is doing better than itself under another strategy. Thus, our results suggest that self-interested agents should adopt the social strategy if social capital is tracked, but without social capital there is no pressure for them to do so.

Recent research has shown that a diverse range of agent-based mechanisms can be effective at managing community energy systems. Allowing agents to form self-organised clusters working to optimise their collective performance has shown to be an effective approach with larger population sizes (Čaušević et al., 2017). There are also more complex algorithms that have the potential to be highly effective at managing decentralised heating systems (Kolen et al., 2017; Dengiz and Jochem, 2020). Our work differs in that it is inherently human facing. A real world implementation of our system could easily operate in a socio-technical manner in which individuals can take over from the virtual agent representing them, setting their own preferences for time-slots and making decisions on whether or not to accept requested exchanges. Utilising a system based on social capital repre-

sented as ‘favours’ would also be easy for the average user to understand facilitating procedural justice.

In conclusion, we have demonstrated a decentralised mechanism for household load balancing that is effective at satisfying agents’ preferences. The benefits of a decentralised mechanism are that it is inherently scalable as more agents are introduced (Petruzzi et al., 2013), and helps to promote privacy and trust by not requiring households to submit their time-slot preferences to a centralised authority. In a real implementation, the exchanges may be performed by software agents running on home gateways installed in households. This could involve various levels of user engagement with the exchange process. The mechanism could also be used alongside differential pricing – households could be given a cheaper rate in their allocated time-slots. Future work should empirically investigate how users perceive distributive and procedural justice both with and without tracking social capital (Powers et al., 2019).

Acknowledgements

This work has been supported by a fellowship to Simon T. Powers from the Keele University Institute of Liberal Arts and Sciences.

References

- Boyd, R. and Richerson, P. J. (1985). *Culture and the Evolutionary Process*. University of Chicago Press, Chicago.
- Čaušević, S., Warnier, M., and Brazier, F. M. (2017). 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)*, pages 568–573. IEEE.
- Dengiz, T. and Jochem, P. (2020). Decentralized optimization approaches for using the load flexibility of electric heating devices. *Energy*, 193:116651.
- Dutta, G. and Mitra, K. (2017). A literature review on dynamic pricing of electricity. *Journal of the Operational Research Society*, 68(10):1131–1145.
- Hollander-Blumoff, R. and Tyler, T. R. (2008). Procedural justice in negotiation: Procedural fairness, outcome acceptance, and integrative potential. *Law & Social Inquiry*, 33(2):473–500.
- Kolen, S., Molitor, C., Wagner, L., and Monti, A. (2017). Two-level agent-based scheduling for a cluster of heating systems. *Sustainable cities and society*, 30:273–281.
- Petruzzi, P. E., Busquets, D., and Pitt, J. (2013). Self organising flexible demand for smart grid. In *2013 IEEE 7th International Conference on Self-Adaptation and Self-Organizing Systems Workshops*, pages 21–22. IEEE.

- Petruzzi, P. E., Busquets, D., and Pitt, J. (2014). Experiments with social capital in multi-agent systems. In Dam, H. K., Pitt, J., Xu, Y., Governatori, G., and Ito, T., editors, *PRIMA 2014: Principles and Practice of Multi-Agent Systems*, Lecture Notes in Computer Science, pages 18–33, Cham. Springer International Publishing.
- Powers, S. T., Meanwell, O., and Cai, Z. (2019). Finding fair negotiation algorithms to reduce peak electricity consumption in micro grids. In *Advances in Practical Applications of Survivable Agents and Multi-Agent Systems: The PAAMS Collection*, Lecture Notes in Computer Science, pages 269–272. Springer International Publishing.
- Putnam, R. D. (1994). Social capital and public affairs. *Bulletin of the American Academy of Arts and Sciences*, 47(8):5–19.
- Rescher, N. (1966). *Distributive Justice: A Constructive Critique of the Utilitarian Theory of Distribution*. The Bobbs-Merrill Company, Inc., Indianapolis.
- Simmons, S. I. and Rowlands, I. H. (2007). TOU rates and vulnerable households: Electricity consumption behavior in a Canadian case study. Technical report, University of Waterloo.
- Stern, P. C., Aronson, E., Darley, J. M., Hill, D. H., Hirst, E., Kempton, W., and Wilbanks, T. J. (1986). The effectiveness of incentives for residential energy conservation. *Evaluation Review*, 10(2):147–176.
- Walker, G. and Devine-Wright, P. (2008). Community renewable energy: What should it mean? *Energy Policy*, 36(2):497–500.