

Demand-Side Management in Smart Grids: an Evolutionary Games Perspective

Daniele Miorandi and Francesco De Pellegrini
CREATE-NET, via alla Cascata 56/D, 38123 - Trento, Italy
{dmiorandi,fdpellegrini}@create-net.org

Abstract—In the context of smart energy grids, demand-side management refers to the ability of dynamically controlling and scheduling energy-consuming tasks. In one potential deployment scenario, smart appliances are controlled by a local intelligent software agents, which implement a given optimization algorithm for scheduling such tasks. The higher the fraction of users adopting such technology, the higher the advantage for the energy operator, due to the ability to control load curve and smooth peaks. At the same time, single users may incur some penalties, related to the fact that energy-consuming tasks may be deferred, thereby causing inconveniences.

In this paper we take a game-theoretical perspective at demand-side management techniques. Tools and solution concepts from evolutionary games are employed: we are interested in the dynamics of the adoption of demand-side management schemes by intelligent software agents. We focus on distributed control schemes that can be enforced by the operator through pricing schemes. Agent-based numerical simulations are provided to validate our theoretical results.

Index Terms—demand side management, customer adoption model, pricing schemes, incentives, game theory, evolutionary stable strategies, replicator dynamics

I. INTRODUCTION

Energy systems are customarily dimensioned based on peak demands. If the demand curve presents a high peak-to-average ratio, this introduces inefficiency and additional costs for the utility providers. Also, when a power supplying company receives more demand for electrical power than it can deliver, it has to resort to rationing of the available electricity to its customers, possibly incurring in contractual violations due to the performed load shedding. Solutions able to smooth peaks in the demand pattern are therefore appealing for operators, in that they can lead to significant savings in both capital and operational expenditures.

In order to smooth peak demands, systems for dynamically adapting demand have been applied in the past to cope especially with “big customers” demands, i.e., big industrial plants with specific type of provisioning, which can be planned “on demand” at the operator side with no need for large scale deployments. However, the situation is recently changing [13], due to the possibility to install smart energy meters on a much larger scale, e.g., covering thousands of domestic appliances in residential areas.

To this respect, the term “demand response” is used to denote a class of new methods for dynamically adapting demands according to some form of control policies. *Demand-side management* (DSM), in particular, refers to a technique

for the automated control and scheduling of the execution of tasks whose completion requires a given amount of energy. Such tasks can be performed, e.g., at the premises of the customer of an electricity provider as in the case of domestic appliances.

The control can be performed in a centralized fashion by the electricity distribution (as in direct load control [12], [3]) or performed in a distributed fashion by intelligent and networked software agents [18], [15]. In the first case, the operator tries to optimize the aggregated demand curve, aiming therefore at maximizing cost savings. In the second case, agents operate on behalf of the user, trying to optimize an objective function that depends on the costs incurred as well as on the delay in the completion of tasks’ execution.

In both cases, the adoption of DSM techniques let the users incur some form of penalty, in the form of costs for upgrading their appliances to DSM-enabled ones and/or of discomfort related to delays in tasks’ execution. The operator shall therefore adopt appropriate pricing and incentive schemes to encourage adoption by customers of DSM technology, while maximizing its own payoff in terms of cost savings. Much literature exists on the investigation of the effect of real-time pricing [4], [12]. Notwithstanding, the question of how to effectively design pricing schemes to foster the adoption of DSM solutions by customers remains unsolved.

In this paper, we present a game-theoretical framework for analysing and predicting the adoption of DSM solutions by users under two different pricing schemes. The pricing schemes analysed differ in the way the costs sustained by a user change as a function of the fraction of customers adopting DSM solutions. In particular, such schemes can be used as transitional solutions in order to incentive the adoption of DSM and to overcome existing barriers.

The most innovative contribution of our work is that we account for interactions among users (or, more precisely: between control agents serving different users). In this way agents can compare their actual electricity costs, and possibly decide to change their strategy to obtain an economical benefit. To model and study this problem, we rely extensively on concepts and tools from evolutionary games theory [23], [28], [10], [6].

We consider a number of pricing schemes that energy distributor may apply and, for each of them, we derive the equilibrium points of the system (in the form of Evolutionary Stable Strategies [28]), i.e., the ones at which, roughly

speaking, the fraction of smart agents adopting DSM schemes does not change over time in the presence of (limited) perturbations, including, e.g., the presence of 'irrational' choices. In particular, we are able to determine various conditions on the pricing schemes that can affect the resulting equilibrium. Not always the adoption of DSM will be able to "invade" the whole customers population: non-trivial equilibrium may arise, in which only a fraction of agents adopt DSM, even in the presence of a fully homogeneous population. Our theoretical findings are validated by means of numerical simulations obtained using an agent-based simulator.

The remainder of the paper is organized as follows. In Sec. II we revise the relevant state-of-art. The system model we use (including the two pricing schemes considered) is elucidated in Sec. III. Results in terms of equilibria and system dynamics are presented in Sec. IV. Validation through comparison with simulation outcomes is provided in Sec. V. Sec. VI concludes the paper pointing out directions for future extensions.

II. BACKGROUND AND RELATED WORK

A. Games, Rationality and Nash Equilibrium

The standard theory of non-cooperative game theory dates back to 1944, with seminal works of von Neumann [26] and later, in the 1950s, of Nash [16]. Since then, game theory has been applied to diverse domains, in particular in economics and in the military strategies field.

According to game theory terminology [8], agents are entities with autonomous decision-making capabilities.¹ Their decision is in terms of strategy to be adopted.

In non-cooperative games, agents cannot or do not want to communicate with other agents and reach a binding agreement on the joint strategy set to be used (e.g., on the way they would share resources). Thus, each agent can react to the current conditions of the system, i.e., a configuration of players' strategies, e.g., a *state*, and try to improve her own utility in a selfish fashion.

A cornerstone of game theory is the notion of Nash equilibria. They are defined as configurations of the strategies of players such that no one of them can increase its payoff by unilaterally deviating from the current profile, provided that all other players do not deviate. The celebrated results of Nash establish the existence of a mixed Nash equilibrium for any finite game, i.e., a game with a finite set of players and strategies and where each user chooses strategies according to a (possibly non-atomic) probability distribution.

Such a notion of equilibrium presents three issues. The first one is on the efficiency of a Nash Equilibrium with respect to a centrally enforced socially optimal strategy profile. In many situations of practical interest, Nash equilibria tend to be rather inefficient in terms of social welfare. The second issue relates to the assumption that players are rational, i.e., that a player

would always choose a best response to the system state, in terms of maximizing her own utility. When agents are humans, they may behave in an irrational way or may not be aware of the complete structure of the game, thereby taking non-optimal decisions. The third issue relates the reachability of an equilibrium: under which conditions the system's dynamics is actually able to drive the system to a certain configuration and how much it would take for the system to stabilize there.

B. Evolutionary Games

The latter two questions have been addressed extensively in the domain of evolutionary games, developed first for biological systems [14], [23]. Evolutionary games deal precisely with the dynamics of strategies within a population, i.e., how many players adopt a certain strategy. The goal is to understand what is the resulting mixture of strategies. E.g., in a two-strategy game, we are interested in the fraction of players adopting the first strategy or the second one. In evolutionary games, assumptions on users' rationality are weak; users are supposed to follow a myopic imitation scheme [21], i.e., to decide on which strategy to adopt based on the immediate advantage they could get. This is done by assuming that players may interact, so that they may revise their strategies based on some form of comparison of the relative utility that they can get using one strategy or the other, and possibly switch to the other one. This may happen for example after a local interaction, i.e., when a player confronts an opponent in a two player game. At a larger scale, the system's evolution is driven by repeated local interactions and by the revision of strategy that the players perform². Overall, evolutionary game theory studies non-cooperative games that are played repeatedly, and identifies distributions of strategies over the population that persist over time because are able to provide the highest fitness of a user.

This type of models – defined initially in the context of biological systems – have been later applied to engineering [24], [2] and economics [21].

1) *Evolutionary Stable Strategies and Replicator Dynamics*: The two main concepts in evolutionary games are the notion of an Evolutionary Stable Strategy (ESS), and the dynamics of the game. The ESS was introduced first by 1972 by the biologist Maynard Smith [14]. ESS provides a notion of stability of a population against invaders, i.e., subsets of players who deviate from a given configuration of the system state. Those invaders are called *mutants*. The presence of an ESS for the system guarantees that whenever the ESS is reached, then the proportions of each strategy in the population do not change over time. Furthermore, at ESS, populations are immune from being invaded by small fractions of mutants.

The second main concept in evolutionary games is that of replicator dynamics. It describes how, in an evolutionary game setting, the distribution of strategies (or 'types') of a population evolve over time. Initially, it was introduced in a biological setting to model how the population size

¹In this paper we will use, with a slight abuse of terminology, the word agents interchangeably for both the entities taking decisions as well as their actual embodiment in a computer program.

²One classical assumption is that interactions occur pairwise and revisions follow a Poisson process: in this case the dynamics can be shown to converge to a replicator dynamics [10].

of different species competing for a given (shared) resource changes over time. In this respect, the payoff of a given agent was meant to represent her reproductive rate. The replicator dynamics is usually based on a set of ordinary differential equations, of the Lotka-Volterra type, that define how the distribution of strategies/types evolves over time, as a function of the difference between the actual payoff of a given strategy and the average one in the population. While all ESSs are asymptotically stable in the replicator dynamics [28], the converse is not true. It can be shown that a strategy profile is an ESS if and only if it represents a strongly stable point for the replicator dynamics [10]. This identifies conditions under which the replicator dynamics converge to ESSs [10], [21]. In the presence of multiple ESSs, the replicator dynamics converges to one of them, depending on the initial conditions. Further, it can be shown that the real dynamics of the system converges, properly scaled, to the replicator dynamics [10], [21]. The replicator dynamics is widely used as a good approximation of the actual dynamics of the system.

It is worth remarking that, while in the biological context the replicator dynamics models the self-organization of the strategy distribution in the population, in the engineering or economic scenario it can model the adoption of certain technology [2] or explain the performance of a system [1], which makes the theory very appealing also to design incentives and pricing schemes able to drive the system to a desired operating point.

One classical framework for evolutionary games starts from the analysis of local interaction made by symmetric two players games, which fits the framework we deal with in the following sections. In such games there are two players using two strategies and the utility attained playing each strategy against any chosen strategy does not depend on the player. A classic yet very illustrative example from evolutionary biology is the *Hawk-Dove game* [23]. In this game, there is a large population of animals, i.e., players. Occasionally two animals find themselves in competition on the same piece of food. An animal can adopt an aggressive behavior or a peaceful one. *Hawk* and *dove* are thus the two alternative strategies that agents (or players) have the option to play in order to gain a valuable resource. The standard evolutionary outcome is a polymorphic population of hawks and doves. Yet, this is not the optimal configuration for all individuals: in fact, when all agents played Dove, the highest reproductive success for all individuals would be attained. The analysis of the Hawk-Dove game provides valuable insight into the conflict between behaviors that are beneficial for the group and individual behavior, and the corresponding population dynamics [28].

2) *Stochastic Evolutionary Dynamics for Games*: Replicator dynamics is a deterministic model. However, evolutionary games are inherently stochastic.

Indeed, depending on the scale one looks at the system, one may need to account for several noise sources. In biological models, for instance, noise is useful to model inherent differences between individuals. A general source of noise is *interaction noise* [6]: individuals in a given population, in fact,

interact at random points in time (following some probability distribution), and this has an impact on the relative density of strategies in the system. Spatial or temporal perturbations in the environment may also be present, leading to the introduction of further elements of randomness. The case of external noise [7], [11] is not considered in this paper

Stochastic evolutionary games [22] is a branch at game theory that looks at the stochastic evolution of the distribution of strategies in an evolutionary game. This is useful in particular for understanding the long-term behaviour of the system.

C. Demand-Side Management and Pricing Schemes

Demand-side management [9] represents a technique for modifying, either directly or indirectly, the aggregated electricity demand curve. Usually, DSM is used to try to smooth peaks in the energy demand curve, allowing electricity operators to make a more effective use of available power plants.

In the recent years, thanks to the introduction to the mass market of smart meters and remotely controlled electric appliances, there has been a surge in research works aimed at designing effective ways to controlling and managing energy demands. Besides direct load control (in which the operator has full control over the customers' appliances), it is imperative to devise schemes to foster the adoption and effective usage by customers of DSM technology, in the form of 'energy boxes' to be installed at the customer's premises and which run smart control agents, acting on behalf of the users. Such schemes can take the form of pricing schemes or incentives. One widely studied approach is to use pricing signals (in particular: real time prices) to let agents adapt autonomously their demand curve in such a way to achieve a suitable aggregated one. In this scheme, agents are assumed to be selfish and willing to minimize their electricity bill, subject to some constraints to ensure usability of their appliances.

Authors of [12] analyze scheduling policies for both interruptible and non-interruptible loads; under causality assumptions the problem that the scheduler needs to solve is a Markov decision process under utility maximization at the provider side side.

In [4], a model for optimally adjusting the hourly load level of a given consumer in response to hourly electricity prices is presented. The objective is to maximise the utility of the consumer subject to a number of constraints on load levels and aggregated consumptions.

The authors of [20] provide an interesting game-theoretical framework (based on algorithmic mechanisms design) for the utility company to effectively price users employing DSM. It is shown that the pricing scheme can be designed in such a way that, by acting selfishly, users end up reaching the optimal system-level operating point.

In [3] a dynamic pricing scheme incentivizing consumers to achieve a suitable aggregate load profile is presented. It is shown that, if users agree to share all their load profiles, an efficient distributed mechanisms (based on cooperative game between consumers) can be introduced.

A number of relevant works has been recently developed within the multi-agent systems community. In [27] the authors propose an agent-based decentralised control mechanism to manage micro-storage in the smart grid based on a dynamic pricing scheme. One work close in spirit to our one is [19], in which the authors study a decentralized demand-side management scheme in which agents coordinate to smooth load peaks. It is shown that through such schemes significant reduction in peak demand can be achieved. In the paper, the authors also introduce an evolutionary game model; they show that under their scheme the usage of demand-side management mechanisms is always convenient. (In game theoretical terms, the usage of DSM is a strictly dominating strategy.) In such conditions the dynamics of the game is straightforwardly converging to the single Nash equilibrium. Our work differs in that we do not make specific assumptions on the DSM algorithm employed, and consider all possible situations arising in terms of user-perceived trade-off between monetary savings and discomfort.

Last, a recent work by Couillet et al. [5] analyses, using mean field games, the impact that the introduction of electric vehicles can have on the variation and dynamics of electricity prices.

III. SYSTEM MODEL

We consider a population of players able to choose between two strategies: i.e., use or do not use DSM techniques. That is, we do assume that software agents deployed at the customers' premises have an embedded DSM algorithm, and can decide whether to use it or not. In case DSM gets used, the user may get a reduction in the energy bill, which depends on (i) the DSM algorithm adopted and (ii) the total number of agents adopting DSM, in that this influences the consumption peaks and thereby the cost exposed by energy operators. At the same time, the user incurs a penalty, due to the fact that energy-consuming tasks may get deferred; again, such penalty depends on the specific DSM algorithm used. The payoff the agent obtains is the negation of the cost she sustains.

The main innovation of our model comes from the fact that we *couple* the behaviour of different agents. That is, we consider models in which the price experienced depends not only on the strategy chosen by a given agents, but also on the *fraction of agents* having opted for the usage of DSM.

A. Pricing Schemes

In this paper we consider two main pricing schemes. They are different in the way prices for DSM adopters (and, hence, costs sustained by the agents adopting DSM), vary as a function of the fraction of the agents' population which adopts the DSM technology. For the sake of simplicity, we assume that agents who do not use DSM techniques experience a cost that is constant, irrespective of the fraction of DSM adopters. (In this way opportunistic, 'freeriding' behaviour get discouraged, and only users adopting DSM are the ones sharing the benefits of their choice.)

- *Reinforcing incentives*: the agents experience a reduction in costs that is monotonically increasing in the fraction of agents adopting DSM. The rationale behind this incentive scheme is that the operator can better control the demand curve (and hence reduce its costs) if there is a larger share of the demand that is performed by DSM-enabled customers and for which the demand can be managed in elastic fashion.
- *Vanishing incentives*: the agents experience a reduction in costs that is monotonically decreasing in the fraction of agents using DSM. While this seems counter-intuitive, it is an approach used in a variety of situations for fostering usage of a new technology with key advantages for the provider (in this case: the electricity distribution operator)³

As we will see in the following, both schemes can perform well as transitional mechanisms for the DSM technology adoption, but attention should be paid to the dimensioning of models' parameters.

B. Notation and Terminology

We denote by $\mathcal{S} = \{0, 1\}$ the set of strategies, where "1" denotes the usage of DSM technology. With standard notation $X = \{(x_1, x_2) \in \mathbb{R}_+^2 : x_1 + x_2 = 1\}$ defines the fraction of the population using strategy $i \in \mathcal{S}$ [21]. In the case at hand a single parameter is sufficient to describe the simplex, x , in order to represent the fraction of agent using DSM: $(x_1, x_2) = (x, 1 - x)$.

The total number of agents is N ; in the following, we will also consider the limiting case in which the population has infinite size.

The population state takes values in the set

$$\mathcal{X}^N = \{(x_1, x_2) \in X : N \cdot x \in \mathbb{Z}^2\}$$

For instance, $N \cdot (x_1, x_2) = (10, 90)$ means that 10 agents do not use DSM and 90 do.

The payoff function associated to playing strategy i when the population is in state x is denoted by $F_i : X \rightarrow \mathbb{R}^N$. We denote by $F_0(x)$ (respectively, $F_1(x)$) the function giving the payoff of a agent deciding to adopt strategy 0 (respectively: 1) when the fraction of the population adopting strategy 1 is given by x .

³An example in the energy sector comes from incentives to the adoption of renewable energy sources (RESs). Let us consider, for the sake of simplicity, the Italian situation. The feed-in scheme (<http://www.gse.it/en/feedintariff/Pages/default.aspx>), first introduced in 2003 with the D.L.387/2003, which is based on the EU directive 2001/77/CE and operational since 2005 provides economic incentives for the deployment and installation of renewable energy sources plants. Over the years the feed-in scheme has gone through a number of programmes (currently the 4th one is active), which have seen the progressive reduction of the incentives provided (approximately 20% reduction from the first to the fourth programme, with incentives dropping progressively by another $\sim 40\%$ until 2016). Such policy was meant to control the increasing deployment of RESs plants. The situation has indeed come to the point at which operators of conventional plants (in Italy mostly thermoelectric ones) are driven by market forces to limit the number of productive hours of their plants. However, below a given limit it becomes uneconomic to operate the plants, whose availability is nonetheless necessary to ensure reliability and security of supply in face of massive intermittent RESs.

In the case of the first pricing scheme (reinforcing incentives), the payoff function takes the following form:

$$\begin{cases} F_0(x) = -\alpha; \\ F_1(x) = -\gamma - \alpha \cdot (1 - \varphi(x)); \end{cases} \quad (1)$$

where $\varphi(\cdot)$ is a monotonically increasing function. To avoid trivial situations, we further assume $\varphi(0) \geq 0$ and $\varphi(1) < 1$.

For the second pricing scheme (vanishing incentives), the payoff function takes the following form:

$$\begin{cases} F_0(x) = -\alpha; \\ F_1(x) = -\gamma - \alpha \cdot (1 - \eta(x)); \end{cases} \quad (2)$$

where $\eta(\cdot)$ is a monotonically decreasing function. To avoid trivial situations, we further assume $\eta(0) < 1$ and $\eta(1) \geq 0$.

Agents may change their strategy according to a revision protocol $\rho: \mathbb{R}^N \times \mathcal{X} \rightarrow \mathbb{R}_+^{N \times N}$. A revision protocol [22] is formally defined as a function that takes as inputs the vector of current payoffs and the population state, and returns non-negative matrices as outputs. The (i, j) -th element of the matrix is the (normalized) rate at which an agent playing strategy i switches to strategy $j \neq i$.

Changes take place according to a Poisson process. Each agent can change her strategy at times $\{t_n(i)\}$, where the sequence is a Poisson process with rate R . At each change opportunity a agent selects an 'opponent', and imitates her only if the opponent's payoff is higher than her own one. In this case, the change takes place with a probability proportional to the payoff difference:

$$\rho_{i,j} = x_j \cdot [F_j - F_i]_+ \quad (3)$$

where $[a]_+ = 0$ if $a \leq 0$ and $[a]_+ = a$ if $a > 0$.

We denote by $Y(t)$ the process describing the population composition (in terms of strategies played) at time t . The process $Y(t)$ takes values in \mathcal{X}^N . Given the assumptions described above, $Y(t)$ is a Markov process, with jumps taking place at times $\{t_n\}_{n \in \mathbb{Z}}$, where the sequence $\{t_n\}$ is a Poisson process with rate $N \cdot R$. We consider the embedded Markov chain $Z_n = Y(t_n^+)$, which has transition probability $\Pi = (\Pi_{x,x+u})$ given by:

$$\Pi_{x,x+u} = \begin{cases} \frac{1}{R} x_i \rho_{i,j}(F(x), x) & \text{if } u = \frac{e_j - e_i}{N}, \\ & i \neq j \in \mathcal{S}; \\ 1 - \sum_{i \in \mathcal{S}} \sum_{j \in \mathcal{S}, j \neq i} \frac{1}{R} x_i \rho_{i,j}(F(x), x) & \text{if } u = 0; \\ 0 & \text{otherwise.} \end{cases}$$

The process $\{Z_n^1\}$ denotes the evolution of the number of agents adopting DSM. It is easy to check that $\{Z_n^1\}$ is a birth-and-death process [21]: in the next section we relate the macroscopic behavior of the system to the agents' payoff induced by reinforcing and vanishing incentives.

IV. EQUILIBRIA AND DYNAMICS

A. Visual Analysis

In order to explain the behaviour of the system, we provide a visual representation of the payoffs experienced by the users—for the two strategies they may choose—as a function

of the fraction of users adopting DSM. As this is meant to provide insight into the resulting system-level behaviour, we will assume for the sake of simplicity that the population size N is infinite, i.e., the variable x can take any value in $[0, 1]$. Extension to the discrete case can be easily handled.

Consider the reinforcing incentive scheme. Let us assume, for the sake of simplicity, that $\varphi(x)$ is linear in x and that $\varphi(0) = 0$. The resulting structure is reported in Fig. 1. There are two options. First, we notice that $F_1(0)$ is always smaller than $F_0(0)$, as $\gamma > 0$ and $\varphi(x) \geq 0$. If the payoff for DSM adopters is always below the dashed line (representing $F_0(x)$), then it is never convenient for users to adopt DSM. This happens if $F_1(1) = -\gamma - \alpha[1 - \varphi(1)] < -\alpha = F_0(1)$. Otherwise, the two payoff functions cross at some point $0 < x_{eq} < 1$. Above this point, it becomes more convenient to adopt DSM. There is therefore a clear dependence on initial conditions. If initially there is a low number of DSM adopters, then it is anti-economical to adopt it. Viceversa, in the presence of a large number of DSM adopters it makes no sense not to adopt it.

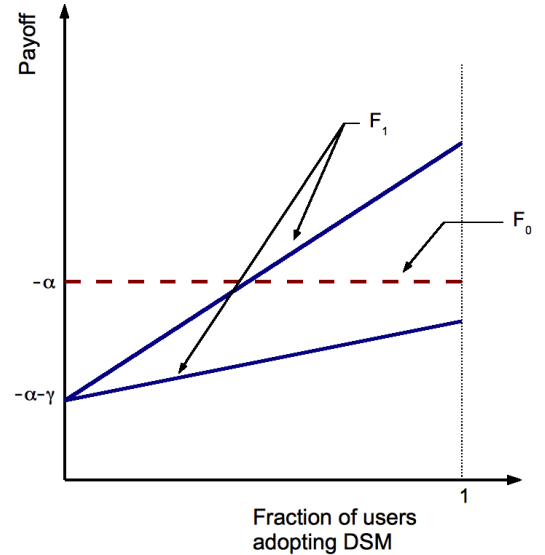


Fig. 1. Payoff as a function of the fraction of DSM adopters for the reinforcing incentives scheme.

Now we move to considering the vanishing incentive scheme. Again, we assume for the sake of simplicity that $\eta(x)$ is linear and $\eta(1) = 0$. We further assume that $\eta(0) > \gamma/\beta$, which ensures that, if nobody is using DSM, there is an incentive for adopting it. The resulting structure is depicted in Fig. 2. Again, there are two cases. If the solid line (corresponding to $F_1(x)$) is always above $F_0(x)$, which happens if $F_1(1) = -\gamma - \alpha[1 - \eta(1)] > -\alpha = F_0(1)$, then it is always convenient to adopt DSM, regardless of the fraction of users choosing it. In the other case, the two curves have a crossing point. In this case, it is convenient to adopt DSM if few users are doing it. But it becomes economically unattractive if there are too many users adopting DSM. At the point where the two curves cross the population

is in equilibrium. Nobody adopting DSM has an advantage in dismantling it, nor nobody not adopting DSM has an incentive to change strategy.




Fig. 2. Payoff as a function of the fraction of DSM adopters for the vanishing incentives scheme.

In the next subsections we will formalize these intuitions using the concepts from game theory described in Sec. II. Some remarks are worth being adding. First, from the reasoning above it is clear that the actual shape of the functions $\varphi(\cdot)$ and $\eta(\cdot)$ has no big impact. What matters is only whether the two curves cross. Similarly, the assumption that $F_0(x) = -\alpha$ could be easily relaxed and replaced by an appropriate monotone function. Again, the system-level behaviour would just depend on whether the two curves $F_0(\cdot)$ and $F_1(\cdot)$ cross.

B. Local interaction

In this section we provide a standard game theoretical view of the evolutionary dynamics. The key result that we will be using is that the ESS set, i.e., the set of evolutionary stable strategies is a subset of the set of Nash equilibria of the local interaction games that occur during interactions between members of the population [21], [25].

In particular, we model the local interaction as a *two player game* with state-dependent payoffs. In the game two users confront each other and compare their relative costs in adopting the DSM technology given the current fraction of agents adopting it (what we call the state). We further assume that the state does not change on the basis of the strategy chosen by the two players (which represents an approximation if the population size is large). The game has the normal form representation [21] reported in Table I. In particular, the case when $\Psi(\cdot) = \phi(\cdot)$ represents the reinforcing incentive scheme, whereas replacing $\Psi(\cdot) = \eta(\cdot)$ we obtain the vanishing incentive scheme.

It is easy to check that if $\gamma < \alpha \cdot \Psi(x)$ the best response against DSM is DSM. Otherwise, it is non-DSM. And vice-versa for playing against non-DSM.

C. Nash Equilibria of Local Interaction

By inspection of the normal form representation, we can determine the possible Nash equilibria of the local interaction games. We need to specialize to the case of the two pricing schemes.

Reinforcing incentives: $\Psi(\cdot) = \phi(\cdot)$: We assume that $\varphi(\cdot)$ is monotone increasing, so that $F_1(\cdot)$ is also increasing. We let $F_0(0) > F_1(0)$; denote x_{eq} the solution of $F_0(x_{eq}) = F_1(x_{eq})$. As the best responses depend not on the which agent plays which strategy but only on how many play strategy '1', we will—with a slight abuse of notation—denote Nash equilibria in terms of the fraction of agents playing strategy '1'. (This corresponds, for mixed equilibria, to a plurality of assignments of strategies over the population.) Then,

Proposition 4.1: Under reinforcing incentives:

- 1) $x = 0$ is always a Nash equilibrium. If $F_1(1) < F_0(1)$, it is unique.
- 2) If $F_1(1) > F_0(1)$ then $x = 1$ is a Nash equilibrium.

Proof: Both cases follow by inspection of the payoff matrix under a continuity argument. For large enough N , since $F_0(0) > F_1(0)$, for any continuous φ it means that there exists $[0, \delta)$, where $0 < 1/N < \delta$ and $F_0(y) > F_1(y)$ in $1 - \delta$.

First, as $F_0(0) > F_1(0)$, there exists an interval $[0, \delta)$, $0 < \delta < 1/N$, such that $F_0(y) > F_1(y)$ for any y , so that $x = 0$ is a Nash equilibrium. If $F_1(1) < F_0(1)$, 0 is a strictly dominating strategy and the uniqueness follows. In case (2), $F_1(1) > F_0(1)$ means that there exists interval $(1 - \delta, 1]$, $0 < \delta < 1$, such that $F_1(y) > F_0(y)$ for any y therein. Again, it follows that in such case $x = 1$ is a Nash equilibrium. The fact that no other Nash equilibria exist follows from a simple monotonicity argument. ■

Vanishing incentives: $\Psi(\cdot) = \eta(\cdot)$: We assume that $\eta(\cdot)$ is monotone decreasing, so that $F_1(\cdot)$ is also decreasing. As before, we let $F_1(0) > F_0(0)$. We have

Proposition 4.2: Under vanishing incentives:

- 1) If $F_1(1) > F_0(1)$, then $x = 1$ is the unique (pure) Nash equilibrium;
- 2) If $F_1(1) < F_0(1)$, then $x = x_{eq}$ is the unique Nash equilibrium.

Proof: The first case can be easily demonstrated by simply noting that '1' is a (strictly) dominating strategy. The second case follows by inspection; showing that x_{eq} is a Nash equilibrium is straightforward, the uniqueness can be demonstrated considering the shapes of $F_1(\cdot)$ and $F_0(\cdot)$. ■

The main difference with the first pricing scheme lies in the second case, in which we have one single mixed equilibrium.

D. ESSs and Stability of the Mean Field Dynamics

The mean field dynamics is characterised by the following differential equation [22], [25]:

$$\dot{x}_i = \sum_{j \in \mathcal{S}} x_j \rho_{j,i}(F(x), x) - x_i \sum_{j \in \mathcal{S}} \rho_{j,i}(F(x), x). \quad (4)$$

	DSM	non-DSM
DSM	$-\gamma - \alpha(1 - \Psi(x)), -\gamma - \alpha(1 - \Psi(x))$	$-\gamma - \alpha \cdot (1 - \Psi(x)), -\alpha$
non-DSM	$-\alpha, -\gamma - \alpha(1 - \Psi(x))$	$-\alpha, -\alpha$

TABLE I

NORMAL FORM REPRESENTATION OF THE LOCAL INTERACTION AS A TWO PLAYERS GAME. $\Psi = \phi$ REPRESENTS THE REINFORCING INCENTIVE SCHEME WHEREAS $\Psi = \eta$ REPRESENTS THE VANISHING INCENTIVE SCHEME.

After few computations it can be shown that, for the revision protocol introduced in (3), (4) becomes the standard replicator dynamics $\dot{x}_i = x_i \cdot \left[F_i(x) - \sum_{j \in \mathcal{S}} x_j F_j(x) \right]$.

We observe that all strict Nash equilibria are by default evolutionary stable strategies [28]. All the Nash equilibria identified above can be shown to be strict and hence also ESSs.

Also, the following result of Hofbauer, Schuster and Sigmund (1979)

Theorem 4.1: Every evolutionary stable strategy is asymptotically stable in the replicator dynamics

Nevertheless, we should also recall that there is a deeper connection of the replicator dynamics (4) with the true dynamics of the system: in particular, there exist results for the convergence of the sample path dynamics to the replicator dynamics [22]. Hence, not only the system state trajectory will converge to an ESS, but also the sample path trajectories of the system will be well approximated by the replicator dynamics equation.

E. Stochastic Dynamics

We now look at the stochastic dynamics of the system. Under the assumptions made above on $\varphi(\cdot)$ and $\eta(\cdot)$, we have:

Proposition 4.3: The following hold:

- 1) For increasing incentives:
 - a) If $F_1(1) < F_0(1)$, the Markov process $\{Z_n^1\}$ is stationary ergodic and converges to a steady state distribution π with $\pi_N = 0$, i.e., no users eventually adopt DSM;
 - b) If $F_1(1) > F_0(1)$ the Markov process $\{Z_n^1\}$ has two recurrent states, $x = 0$ and $x = 1$; all other states are transient;
- 2) For vanishing incentives:
 - a) If $F_1(1) > F_0(1)$, the Markov process $\{Z_n^1\}$ is stationary ergodic and converges to a steady state distribution π with $\pi_N = 1$, i.e., all users eventually adopt DSM;
 - b) If $F_1(1) < F_0(1)$, the process $\{Z_n^1\}$ is irreducible and thus $\pi_i > 0 \forall i$. The limiting distribution can be numerically computed using local balance conditions.

Proof: We deal separately with the four cases above: (1.a) Regardless of the state Z_n^1 , the process can only move to lower value, as in any case it is convenient to drop the usage of DSM. The process is therefore a pure death one and the result follows straightforwardly.

(1.b) First, we can check by inspection that $x = 0$ and $x = 1$ are absorbing states. By continuity there exist a value $\tilde{x} \in (0, 1)$ such that $F_1(\tilde{x}) \leq F_0(\tilde{x})$ and $F_1(\tilde{x} + \frac{1}{N}) > F_0(\tilde{x} + \frac{1}{N})$. Given the expression for the revision protocol used, we can easily check that for $x < \tilde{x}$ the process has negative drift. Similarly, it can be checked that the drift is positive for $x > \tilde{x}$, and hence the result follows. In this case the value at which Z_n^1 settles depends on the initial condition.

(2.a) Regardless of the state Z_n^1 , the process can only move to higher values, as in any case it is convenient to adopt DSM. The process is therefore a pure birth one and the result follows straightforwardly.

(2.b) By inspection it can be checked that the chain has no absorbing states. As the process Z_n^1 takes finite values, the result follows straightforwardly. ■

V. NUMERICAL VALIDATION

In order to validate the findings presented in the previous section, we implemented the proposed models in a Java-based agent simulator. (In order to allow independent experts to repeat the experiments, the code developed has been released under an open source license and can be found at: <http://dl.dropbox.com/u/15242937/e-EnergyJavaCode.zip>). Each agent corresponds to a user. Agents interact randomly and can change their strategy according to the revision protocol defined in Sec. III. We used $N = 100$ as population size. Each simulation run corresponded to 2000 interactions among agents, intervals between subsequent interactions being modelled as an exponentially distributed random variable with average 100 s.⁴

The first set of experiments was focused on the reinforcing incentives scheme. We first studied the convergence of the system dynamics, in terms of fraction of agents using DSM techniques. We considered $\varphi(x) = \beta x^2$, $\gamma = 1.5$ and $\alpha = 5$. In these cases the two curves $F_0(x)$ and $F_1(x)$ cross at $x^* = \sqrt{\frac{\gamma}{\alpha\beta}}$. We first considered an initial fraction of DSM adopters equal to 30% and run simulations for $\beta \in \{0.6, 1, 2\}$. According to the results in the previous section, it is easy to check that in these cases the dynamics should converge to the pure strategy “do not adopt DSM”. The results, for three sample paths, are reported in Fig. 3. The arrow indicates growing values of β . As it can be seen, in all cases the users settle to not adopting DSM, regardless of the value of β , as predicted by the model. The value of β only changes the speed

⁴This value is purely indicative, as changes in it would just imply a rescaling of the time axis, but not a qualitative change in the resulting dynamics.

at which convergence takes place.

We then considered a different, much larger value of initial DSM adopters, i.e., 75%. With the same set of parameters described above, in this case the results in Sec. IV predicts convergence to the pure strategy “adopt DSM”. The resulting dynamics is depicted in Fig. 4. Also in this case the prediction of the models are validated, and the value of β defines only the speed at which the convergence takes place.

To get a more comprehensive picture, we fixed $\beta = 1$ and varied the initial fraction of DSM adopters. The results, averaged over 10 runs, with error bars for the minimum and maximum value over a single run, are reported in Fig. 5. As it can be seen, there is a sharp phase transitions around $\beta \approx 0.55$. Again, this is consistent with the analysis carried out in the previous Section. The large difference between runs at $x_1(0) = 0.55$ indicates that the system is in proximity of a tipping point, with the stochastic effects related to the interactions among agents can drift the system towards one of the two ESSs.




Fig. 3. Sample paths of the dynamics of the fraction of users adopting DSM, reinforcing incentives, $\varphi(x) = \beta x^2$, $x_1(0) = 0.3$, $\gamma = 1.5$, $\alpha = 5$, $\beta \in \{0.6, 1, 2\}$. The arrows indicates increasing values of β .




Fig. 4. Sample paths of the dynamics of the fraction of users adopting DSM, reinforcing incentives, $\varphi(x) = \beta x^2$, $x_1(0) = 0.75$, $\gamma = 1.5$, $\alpha = 5$, $\beta \in \{0.6, 1, 2\}$. The arrows indicates increasing values of β .

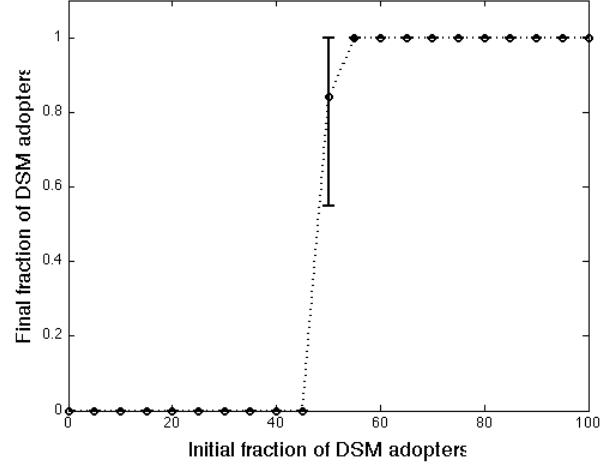


Fig. 5. Final fraction of users adopting DSM vs. initial fraction of DSM adopters, reinforcing incentives, $\varphi(x) = \beta x^2$, $\gamma = 1.5$, $\alpha = 5$, $\beta = 1$.

The second set of experiments was focused on the vanishing incentives scheme. In this case we considered a linear incentive, i.e., $\eta(x) = \frac{1-x}{2}$. We took $\alpha = 1.5$ and $\gamma = \beta$. The crossing point (computed assuming an infinite population model) is given by $x^* = 0.4$. Given the results in Sec. IV on the presence of one single ESS and the convergence of the mean field dynamics, we do expect convergence to x^* regardless of the initial conditions. We run experiments varying the initial fraction of DSM adopters. Results are plotted, in terms of sample paths for $x_1(0) \in \{0.1, 0.45, 0.9\}$ in Fig. 6. It can be seen that the behaviour of the system is consistent with the predictions offered by the model.




Fig. 6. Sample paths of the dynamics of the fraction of users adopting DSM, vanishing incentives, $\eta(x) = \frac{1-x}{2}$, $x_1(0) = 0.1, 0.45, 0.9$, $\gamma = 1.5$, $\alpha = 5$.

A. DSM: a Reality Check

In the previous sections, in line with the spirit of this paper, we made no specific assumption on the behaviour and

performance of the DSM scheme to be applied. In this section we aim at presenting a reality check, in order to understand the actual savings that a customer can obtain by means of DSM in the presence of real-time prices and assess whether real-time prices can actually provide sufficient incentives to adopt DSM.

Modelling the real impact and behaviour of DSM schemes with a sufficient level of generality represents a challenging issue due to four main reasons:

- The load profile varies heavily according to a number of contextual factors, including: period of the year, weather, location, family composition, working habits etc.
- A huge variety of DSM mechanisms has been proposed in the literature, including different architectures, algorithms and signalling schemes. Even a minor difference in the mechanism can have a rather large impact on the resulting performance.
- In different countries different pricing schemes are applied by distributors. Further, energy prices vary significantly from day to day due to a number of contextual factors (including weather, season etc.).
- It is difficult to estimate the 'discomfort' caused by the adoption of DSM schemes. In general, the more a job is deferred, the higher the discomfort for the user. However, it is very difficult to estimate it reasonably, as this depends on a number of factors (related to the type of job but also to cultural aspects) for which no reference model is currently available.

We made a number of simplifying assumptions. As for prices, we used data from Ameren (IL), corresponding to the real-time prices (RTP) for Zone I measured on 30/4/2012⁵. The corresponding prices (in \$/ kWh) are reported in Fig. 7.




Fig. 7. Real-time prices from Ameren measured on April 30th, 2012.

We also used data and models from [17] in terms of 'typical' load profile for households (data for week days in Table B.II

therein). Appliances are grouped in ten categories: (i) stove and oven (ii) microwave and coffee maker (iii) refrigerator and freezer (iv) dishwasher (v) clothes-washer and tumble dryer (vi) television and video recorder (vii) radio/player (viii) PC and printer (ix) lighting (x) other loads. The profile is divided into twenty-four one-hour intervals. Each category is assigned a 'priority level', with the starting time of jobs from category (iii) being possibly deferred up to one hour, and those of categories (iv) and (v) being possibly deferred up to six hours. The loads from the other categories are considered non deferrable.

We considered two DSM schemes. In the first one ('Simple DSM'), a simple randomization scheme is used, with the starting time of a deferrable job being deferred by a random quantity drawn from a uniform distribution (within the boundaries expressed by the maximum deferral period). It is worth remarking that in such scheme no real-time pricing is used, but a simple randomization local scheme can be used. While this scheme is simplistic, it has the advantage of not requiring communication between the local control agent and the operator.⁶

In the second one ('Look-ahead DSM') we assumed the presence of an 'oracle' being aware of the prices as they will evolve in the future hours. Accordingly, jobs are deferred to the time corresponding to the minimum electricity price (within the boundaries expressed by the maximum deferral period). This scheme, which is clearly idealistic, provides an upper bound on the savings obtained by any implementable scheme.

The resulting profiles are reported in 8. A simple numerical calculation reported that for the Simple DSM scheme the peak electricity load reduction was below 1%, with an expected savings for customers of 2.06%. With the Look-Ahead DSM scheme the peak electricity load was increased by more than 50%, with an expected savings for customers of 8.25%.

These figures were obtained assuming that the prices are independent from the fraction of users adopting DSM. The following conclusions can be drawn:

- The actual savings a single user can get from DSM with real-time prices are marginal and unlikely to foster by themselves a wide adoption.
- The usage of real-time prices may lead to the creation of peaks in the residential users' profile. A careful design is required to ensure electricity operators can actually benefit from the widespread adoption of DSM solutions.

VI. CONCLUSIONS

In this paper we have analysed the effect of two different pricing schemes on the adoption by intelligent software agents of demand-side management techniques. In one of the pricing schemes, the incentive for agents to adopt DSM is increasing in the fraction of agents using it while in the second one it is decreasing. Leveraging on concepts and results from evolutionary game theory, we have shown that in the first

⁶A DSM scheme based on real-time pricing is likely to achieve much higher performance gain at the cost of increased complexity.

⁵<https://www2.ameren.com/RetailEnergy/realtimeprices.aspx>

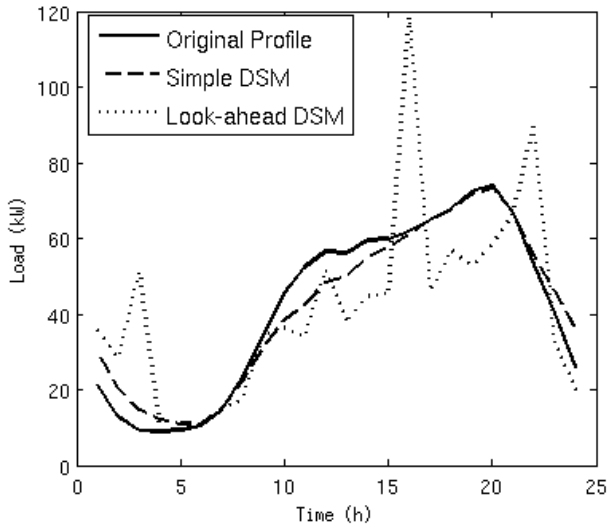


Fig. 8. Resulting load profiles (original, Simple DSM and Look-ahead DSM).

case—besides trivial cases—there exists a phase transition in the eventual adoption of DSM. In such a case, an electricity distribution operator should provide means for ensuring that a sufficiently large amount of agents adopt DSM before the incentives kick in. In the second case, at the opposite, the system converges (regardless the initial distribution of DSM adopters) towards a mixed equilibrium, in which only a fraction of the agents' population uses DSM. The theoretical findings have been validated by comparison with the outcomes of agent-based simulations.

The present work represents a first step towards a coherent framework for designing pricing and incentive schemes for DSM adoption. Future extensions will include heterogeneous population (with different users having different consumption requests/profiles) and enhance the level or realism by considering models for the users' discomfort caused by DSM.

REFERENCES

- [1] ALTMAN, E., AZOUZI, R. E., HAYEL, Y., AND TEMBINE, H. Evolutionary power control games in wireless networks. In *Proc. of IFIP Networking* (2008), pp. 930–942.
- [2] ALTMAN, R. ELAZOUZI, Y. H., AND TEMBINE, H. The evolution of transport protocols: An evolutionary game perspective. *Elsevier Computer Networks Journal* 53, 10 (2009), 1751–1759.
- [3] CARON, S., AND KESIDIS, G. Incentive-based energy consumption scheduling algorithms for the smart grid. In *Proc. of IEEE SmartGridComm* (Gaithersburg, October 2010).
- [4] CONEJO, A., MORALES, J., AND BARINGO, L. Real-time demand response model. *IEEE Trans. On Smart Grid* 1, 3 (Dec. 2010), 236–242.
- [5] COUILLET, R., PERLAZA, S., TEMBINE, H., AND DEBBAH, M. A mean field game analysis of electric vehicles in the smart grid. In *Computer Communications Workshops (INFOCOM WKSHPS), 2012 IEEE Conference on* (March 2012), pp. 79–84.
- [6] FREY, E. Evolutionary game theory: Theoretical concepts and applications to microbial communities. *Physica* 389 (2010), 4265–4298.
- [7] FUDENBERG, D., AND HARRIS, C. Evolutionary dynamics with aggregate shocks. *Journal of Economic Theory* 57 (1992), 420–441.
- [8] FUDENBERG, D., AND TIROLE, J. *Game theory*. MIT Press, Cambridge, MA, 1991.
- [9] GELLINGS, C., AND CHAMBERLIN, J. *Demand-side management: Concepts and methods*. The Fairmont Press Inc., Lilburn, GA, 1987.
- [10] HOFBAUER, J., AND SIGMUND, K. Evolutionary game dynamics. *American Mathematical Society* 40, 4 (2003), 479–519.
- [11] IMHOF, L. The long-run behavior of the stochastic replicator dynamics. *The Annals of Applied Probability* 15 (2005), 1019–1045.
- [12] KIM, T., AND POOR, H. Scheduling power consumption with price uncertainty. *IEEE Trans. On Smart Grid* 2, 3 (Sept 2011), 519–527.
- [13] KRISHNA, R. Meters of tomorrow. *IEEE Power and Energy Magazine* (March 2008), 92–94.
- [14] MAYNARD SMITH, J. *Game theory and the evolution of fighting—*. Cambridge University Press, Cambridge, UK, 1972.
- [15] MCARTHUR, S., DAVIDSON, E., CATTERSON, V., DIMEAS, A., HATZIARGYRIOU, N., PONCI, F., AND FUNABASHI, T. Multi-agent systems for power engineering applications - part I: Concepts, approaches, and technical challenges. *IEEE Trans. On Power Systems* 22, 4 (2007), 1743–1752.
- [16] NASH, J. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences of the United States of America* 36, 1 (1950), 48–49.
- [17] PAATERO, J. V., AND LUND, P. D. A model for generating household electricity load profiles. *International Journal of Energy Research* 30, 5 (2006), 273–290.
- [18] PIPATTANASOMPORN, M., FERROZE, H., AND RAHMAN, S. Multi-agent systems in a distributed smart grid: Design and implementation. In *Power Systems Conference and Exposition, 2009. PSCE '09. IEEE/PES* (March 2009), pp. 1–8.
- [19] RAMCHURN, S., VYTELINGUM, P., ROGERS, A., AND JENNINGS, N. Agent-based control for decentralised demand side management in the smart grid. In *Proc. of AAMAS* (Taipei, Taiwan, 2011), pp. 5–12.
- [20] SAMADI, P., SCHÖBER, R., AND WONG, V. Optimal energy consumption scheduling using mechanism design for the future smart grid. In *Proc. of IEEE SmartGridComm* (Brussels, Belgium, Oct. 2011).
- [21] SANDHOLM, W. H. *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge, MA, 2010.
- [22] SANDHOLM, W. H. Stochastic evolutionary game dynamics: Foundations, deterministic approximation, and equilibrium selection. In *Evolutionary Game Dynamics*, K. Sigmund, Ed. American Mathematical Society, 2011, pp. 111–141.
- [23] SMITH, M. *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, UK, 1982.
- [24] TEMBINE, H., ALTMAN, E., AZOUZI, R. E., AND HAYEL, Y. Evolutionary games in wireless networks. *IEEE Transactions on Systems, Man, and Cybernetics, Part B* 40, 3 (2010), 634–646.
- [25] VINCENT, T. L., AND BROWN, J. S. *Evolutionary Game Theory, Natural Selection, and Darwinian Dynamics*. Cambridge University Press, 2005.
- [26] VON NEUMANN, J., AND MORGENSTERN, O. *Theory of games and economic behavior*. Princeton University Press, 1944.
- [27] VYTELINGUM, P., VOICE, T., RAMCHURN, S., ROGERS, A., AND JENNINGS, N. Theoretical and practical foundations of agent-based micro-storage in the smart grid. *Journal of Artificial Intelligence Research* 42 (2011), 765–813.
- [28] WEIBULL, J. *Evolutionary Game Theory*. MIT Press, Cambridge, MA, 1997.