

Autonomous Agents in Future Energy Markets: The 2012 Power Trading Agent Competition

Wolfgang Ketter and **Markus Peters**
Erasmus University Rotterdam
{*wketter,peters*}@rsm.nl

John Collins
University of Minnesota
jcollins@cs.umn.edu

Abstract

Sustainable energy systems of the future will need more than efficient, clean, and low-cost energy sources. They will also need efficient price signals that motivate sustainable energy consumption behaviors and a tight real-time alignment of energy demand with supply from renewable and traditional sources. The Power Trading Agent Competition (Power TAC) is a rich, competitive, open-source simulation platform for future retail power markets built on real-world data and state-of-the-art customer models. Its purpose is to help researchers understand the dynamics of customer and retailer decision-making as well as the robustness of proposed market designs. Power TAC invites researchers to develop autonomous electricity broker agents and to pit them against best-in-class strategies in global competitions, the first of which will be held at AAAI 2013. Power TAC competitions provide compelling, actionable information for policy makers and industry leaders. We describe the competition scenario, demonstrate the realism of the Power TAC platform, and analyze key characteristics of successful brokers in one of our 2012 pilot competitions between seven research groups from five different countries.

Introduction

Energy markets worldwide are undergoing momentous change. In the European Union alone, required investments into the liberalized, sustainable, smart electric grid of the future are projected to run up to more than two trillion Euro between 2011 and 2050, with total energy systems cost rising from 10.5% to 14.6% of the continent's GDP (European Commission 2011). Key drivers behind these developments are the political desire for improved economic welfare among consumers, as well as an increased share of electricity production from renewable sources. Governments worldwide are adopting ambitious agendas to promote these goals, and research into "secure, clean and efficient energy" has been identified as one of the key areas that require the immediate attention of the scientific community (European Commission 2013). Core issues to be addressed include (a) the need for decentralized control mechanisms that deliver the same degree of reliability previously afforded by monopolistic providers, and (b) the need for novel incentives for

customers to shift electricity usage to times when renewable sources are available (Ramchurn et al. 2012).

At their core, these are problems in decentralized, real-time economic decision making that have long been under study in the autonomous agents community, e.g. (Wellman, Greenwald, and Stone 2007), and several authors have recently used the results of this work to design agents that solve selected aspects of the issues outlined above, e.g. (Peters et al. 2013; Reddy and Veloso 2011; Vytelingum et al. 2010). While the results of these authors are promising, we find that they are limited in two important ways:

Limited Scope Energy markets are based on a complex interplay among several markets on which future obligations with intricate properties are traded. Few research groups command the resources required to build a robust, validated simulation of this demanding environment before embarking on their actual research.

Limited Competitiveness and Comparability Each study starts from a limited, self-built environmental model, making results difficult to compare. This limits the impact of research results and reduces the incentives for researchers to be involved in this work.

We address these limitations with the **Power Trading Agent Competition** (Power TAC, Ketter et al. 2012b, see also www.powertac.org). Power TAC is a rich, open-source simulation platform for future retail power markets coupled with a series of annual competitions that challenge participants to build autonomous, self-interested agents that compete directly with each other in this demanding environment. We are hosting the first official Power TAC competition at AAAI 2013 in Bellevue, WA.

Power TAC advances the state of the art in five important ways: (1) It is the first comprehensive simulation platform for future retail power markets. It supports research into mechanism design, autonomous retail and wholesale electricity trading, and intelligent automation techniques centered on human preferences. (2) It provides a standardized research infrastructure, alleviating the need for costly up-front creation of environmental models, lowering barriers to entry for new researchers, and promoting comparability between scientific studies in the field. (3) Its competitive nature and the availability of state of the art benchmark strategies will foster innovation. (4) The platform is used and

supported by a growing community of researchers and developers who contribute state-of-the-art models for all facets of the environment, leading to continuous improvements in the simulation's realism and sophistication. And (5) Power TAC advances beyond earlier Trading Agent Competitions by providing extensive facilities for experiment management, data extraction, data visualization and analysis, and mixed-mode games with human participation.

The paper is organized as follows. First, we give an overview of related work in the agent, simulation and energy market literatures. Next, we describe the competition scenario in some detail. And finally, we present a selection of interesting insights based on one of our 2012 pilot tournaments.

Related Work

Electricity markets are undergoing a transition from regulated monopolies to decentralized markets (Joskow 2008), but so far the retailers in these markets are almost entirely limited to purchasing power in the wholesale markets and delivering it to their customers (Fleten and Pettersen 2005); they have not had to deal with significant volumes of power production among their customers. A critical unanswered question is the extent to which self-interested behaviors of market participants can effectively supplement hierarchical control of the physical infrastructure to balance supply and demand in such an environment. To answer this important question we base our work on two major research streams, **computational sustainability** and **competitive agent-based simulations**.

The newly emerging computational sustainability field studies the application of AI and other computational techniques to sustainability issues, and smart grids are one focus area for that community. For example, Voice et al. (2011) explore market-based strategies for controlling the use of "home" and other micro energy storage. And Ermon et al. (2012) formalize the problem of optimizing real-time energy management of multi-battery systems as a stochastic planning problem, and propose a novel solution based on a combination of optimization, machine learning and data-mining techniques.

Among the many important open questions and research challenges posed by a power grid with large numbers of active participants, e.g. (Ramchurn et al. 2012), only few can be addressed by straightforward game-theoretic analysis (de Weerd, Ketter, and Collins 2011). The others are sufficiently complex that they cannot be effectively addressed by formal methods. To address these more complex issues, a simulation-based technique known as Agent-based Computational Economics (ACE, Tesfatsion 2002) has been used to study electrical wholesale power markets, for example (Sun and Tesfatsion 2007; Reddy and Veloso 2011; Peters et al. 2013). Like other Trading Agent Competition scenarios (Ketter and Symeonidis 2012), Power TAC extends the ACE paradigm by creating a rich economic simulation and inviting research teams to develop their own software agents to play the role of power retailers in the simulation, and to enter them in annual competitions.

The Power TAC Scenario

The main elements of the Power TAC scenario are shown in Figure 1. The scenario models a competitive retail power market in a medium-sized city, in which consumers and small-scale producers of electricity may choose among a set of alternative electricity power providers, represented by competing broker agents. The brokers are self-interested, autonomous agents, built by individual research groups to participate in the competition; the remainder of the scenario is modeled by the Power TAC simulation platform. In the real world, brokers could be energy retailers, commercial or municipal utilities, or cooperatives.

In a tournament environment, simulations are run with different numbers and combinations of broker agents, and the agent that is most profitable over a range of scenarios is the winner. Importantly, profit maximization does not preclude other social desiderata such as fairness, utilization of renewables, or certain levels of electric vehicle market penetration. By properly setting the market's **economic mechanisms** (Dash, Jennings, and Parkes 2003), market designers can create incentive structures that lead self-interested, profit-maximizing brokers towards socially desirable outcomes. In a research environment, Power TAC may be extended with experimental user-implemented mechanisms that can then be subjected to broker competition to determine their effectiveness.

The simulation proceeds in a series of discrete timeslots, each representing one hour in the simulation world. A typical tournament simulation runs approximately 60 days of simulated time, although much longer simulations are possible. Time advances by one timeslot every five seconds, so a simulation completes in about two hours. The actual duration of the scenario is stochastic, to minimize the opportunity for brokers to exploit a predictable "end-of-game" situation that, while it might win tournaments, has little research value or relationship to the real world.

Customers and Tariff Market Brokers interact through a retail tariff market with customer models that simulate the households and businesses of a small city. Some customers are equipped with solar panels and windmills, producing as well as consuming power. All customers are assumed to be equipped with smart meters; consumption and production is reported every hour. Many customer models also include **controllable capacities** or demand-side management capabilities such as heat pumps or water heaters. These devices can be remotely enabled or disabled to offset imbalances or control costs, in exchange for lower rates.

Customer models exhibit sensitivity to weather conditions (real-world weather reports and forecasts are used) and calendar factors such as day of week and hour of day. The models also respond to price changes (Gottwalt et al. 2011) and have a range of preferences over tariff terms. For example, some are willing to subscribe to variable-price tariffs if they have the opportunity to save by adjusting their power usage, while others are willing to pay higher prices for the simplicity of fixed-rate or very simple time-of-use tariffs. This behavior is supported by real-world pilots with dynamic pricing schemes (EPRI 2012).

Many customer models are capable of adaptive capacity

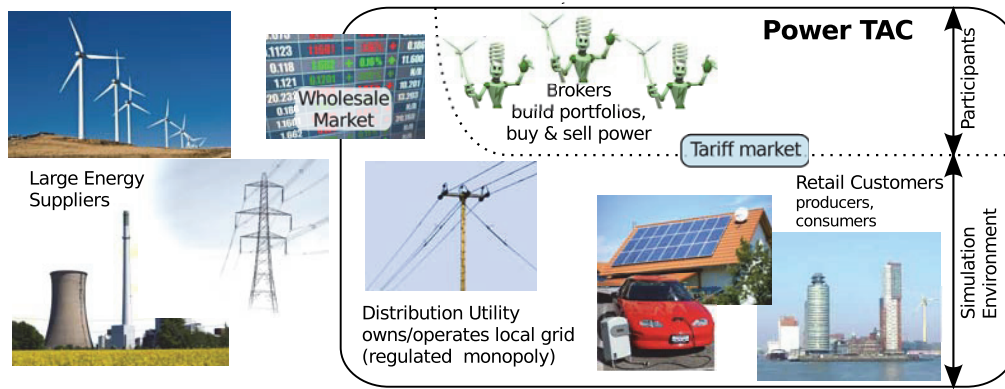


Figure 1: Main elements of the Power TAC scenario. Brokers are the competitors; markets, customers, energy suppliers, and distribution utility are modeled by the Power TAC simulation platform.

management, allowing them to evaluate various possibilities for capacity shifting and choose amongst the ones best suited to the applicable tariff rates, while also considering the potential choices of other customers that may be on the same tariff (Reddy and Veloso 2012).

Tariff contracts may include usage-based and per-day charges, fixed and varying prices for both consumption and production of energy, rates that apply only above a specified usage threshold, signup bonuses, and early-withdrawal penalties. Separate contracts may be offered for charging electric vehicles, which could limit charging during high-demand periods, or even offer to pay customers for feeding energy back into the grid at certain times. Variable prices may follow a fixed schedule (day/night pricing, for example), or they may be fully dynamic, possibly with specified advance notice of price changes.

Wholesale Market and Generating Companies Brokers may buy and sell energy from retail customers, and they may buy and sell energy for future delivery in a wholesale market, which is modeled on real-world markets such as the European and North American day-ahead wholesale power markets. At any given time, brokers may place orders in the wholesale market to buy or sell power in 24 separate auctions, the first for delivery in the following timeslot, and the last 24 hours in the future. On the supply side, the Power TAC platform includes large producers (or **Gencos**) that simulate utility-scale power suppliers who sell their output through the wholesale market. These suppliers represent different price points and lead-time requirements, e.g. fossil and nuclear power plants, gas turbines, and wind parks.

Distribution Utility The Distribution Utility (DU) models a regulated monopoly or government entity that owns and operates the physical facilities (distribution lines, transformers, etc.) and is responsible for real-time balancing of supply and demand within the distribution network. It does this primarily by operating in the **regulating market**, the real-time facet of the wholesale market, and by exercising demand and supply controls provided by brokers. The associated costs are allocated to the brokers responsible for the imbalance. In the real world, this balancing responsibility is typically han-

dled higher in the grid hierarchy, by the Independent System Operator (ISO, North America) or Transmission System Operator (TSO, Europe); the simulation implements a generalization of proposals to move some balancing responsibility to the distribution level (Strbac 2008).

Power TAC is a market simulation. It abstracts away most technical considerations, such as power factor and distribution losses, that arise from non-ideal behaviors of the real power infrastructure. As long as volumes of delivered power do not approach physical constraints, we can treat most such phenomena as discounts, and fold them into the market structures. For example, distribution losses can be roughly accounted for by charging brokers a per-kilowatt-hour fee for delivering power to their customers.

Brokers Brokers develop customer portfolios by offering tariff contracts to a population of anonymous residential and business customers, and by negotiating individual contracts with larger customers (such as major manufacturing facilities, or greenhouse complexes with many Combined Heat and Power (CHP) units).¹ Because controllable capacities can reduce costs significantly, brokers can offer special tariffs for them, and then make offers to the DU for the right to exercise them to reduce imbalances. Given a portfolio of customers, brokers compete with each other in the wholesale market to minimize the cost of power they deliver to their consuming customers, and to maximize the value of power delivered to them by their producing customers.

Insights from the 2012 Pilot Competitions

We have hosted several pilot competitions on the Power TAC platform, including competitions at IJCAI 2011 in Barcelona, at AAMAS 2012 in Valencia, and at IEEE SGTEP 2012 in Nuremberg.² Teams from Croatia, Germany, Greece, Korea, Mexico, Netherlands, UK, and the US entered brokers in these tournaments, providing us with a full-

¹The negotiation feature was not implemented for the 2012 tournaments.

²International Joint Conference on Artificial Intelligence, Autonomous Agents and Multi-Agent Systems, and IEEE Conference on Smart Grid Technology, Economics, and Policies, respectively

Broker	Research Group
AstonTAC	Aston University Birmingham
CrocodileAgent	University of Zagreb
LARGEpower	Erasmus University Rotterdam
Mertacor	Aristotle University Thessaloniki
MinerTA	University of Texas at El Paso
SotonPower	University of Southampton
UTest	University of Texas at Austin

Table 1: Participants in the 2012 Nuremberg pilot competition

scale validation of Power TAC, and broker developers with opportunities to evaluate their strategies.

A credible simulation must emulate interesting real-world phenomena, while abstracting away details that are unlikely to contribute to the phenomena of interest. As part of our validation, we expected several characteristic phenomena to emerge from interaction between brokers and the Power TAC platform. These included:

Competitiveness Broker performance should depend on broker strategies, not on random draws. We expected consistent, discernible performance differences between high and low performers in the field.

Resilience to Competition We expected the best strategies to keep their competitive edge under competition from additional brokers.

Balancedness We expected that the best strategies would perform competitively in all three markets: retail, wholesale, and regulating market.

Economic Realism We expected to re-discover a number of fundamental economic truths in the competition results: prices should correlate with demand, wholesale prices should be higher for shorter leadtimes, and profit margins should fall with increasing competition.

We analyzed 51 games from the Nuremberg 2012 pilot competition.³ Our analyses give insights into the key success factors behind high-performing broker strategies, and suggest important areas of future work for broker developers. Table 1 lists the participating brokers.

To get a sense of overall performance, we first analyzed broker *profit shares* – the percentage of each broker’s profit in all profits made during a game (see Figure 2). Interestingly, we find pronounced differences in terms of profit share *magnitude*, *certainty*, and *resilience to competition*. AstonTAC, for example, does well on all three counts, with profits that are high on average, stable over different games, and relatively unaffected by the additional competition introduced in 5-player games. Mertacor and LARGEpower, on the other hand, play strategies that perform well in some cases, but that are not as consistent. To understand the reasons for these differences we analyzed the brokers’ actions in Power TAC’s three principal markets: retail, wholesale, and regulating market.

³The Nuremberg competition used the most recent beta version of the Power TAC platform; for this analysis, we excluded games




Figure 2: Average profit share of brokers in the Nuremberg 2012 pilot competition. Certainty is computed as the inverse sample standard deviation. Small numbers indicate the number of games on which a data point is based. The default broker is an unsophisticated strategy provided by Power TAC itself and models the incumbent monopoly.

Figure 3 shows the strategic positions of brokers in the retail market for electricity consumption.⁴ As expected, there is a clear relationship between the price level offered on electricity tariffs and the share of customer demand served. Mertacor, Aston TAC, and UTest are high volume players, offering competitive rates in exchange for significant shares of the retail market. Mertacor, in particular, consistently captures the majority of the market at extremely low prices which cut into its own margins (see Figure 2). SotonPower and AstonTAC are interesting in that their successes appear to hinge on fundamentally different causes: while Aston TAC aims to attract large volumes in the retail market, SotonPower’s success is apparently based on a more moderated combination of larger markups and other factors discussed below. Importantly, the retail rates offered are typically less than half the cost of the rates offered by the incumbent monopoly (default broker), which illustrates the benefits of competitive retail pricing to customers. Somewhat surprisingly, we find no discernible variation in retail markups over different game sizes. This suggests that the early-stage brokers studied here are not yet fully exploiting their strategic options.

To offset the (expected) consumption of their retail customers, brokers must acquire future supply commitments in

that were affected by technical flaws of the beta.

⁴Only AstonTAC and UTest also *purchased* power in the retail market during the games that we analyzed; volumes were small in relation to wholesale market activities.

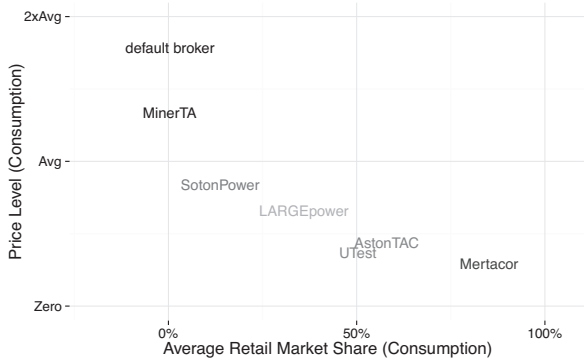


Figure 3: Retail market position of participating strategies. Lighter coloring indicates higher market share uncertainty.

the wholesale market. Figure 4 illustrates brokers' strategic positions in this market. SotonPower's low-volume strategy, likely combined with high forecasting accuracy, allows it to cover its electricity needs at extremely low prices, which explains at least part of its overall success. The relatively weak brokers LARGEpower, Mertacor, and UTest are, conversely, paying prices that are up to 1.6 times above market average. Note that this is *not* solely a consequence of large volumes and the corresponding need to invoke marginal, high-cost providers of electricity. AstonTAC, for example, acquires larger amounts of electricity than LARGEpower at lower prices on average, suggesting better timing of wholesale purchases or smarter use of bid prices. Mertacor pays the most for wholesale power; together with its low-cost tariffs, this largely explains its relatively weak performance.




Figure 4: Wholesale market position of participating strategies. Lighter coloring indicates higher market share uncertainty.

We next looked at the market for balancing power (regulating market, see Figure 5). This market is implicit in that residual imbalance between consumption and production is automatically offset by the DU in real-time. The penalties (or rewards) that a broker incurs in the process are, however, dependent on the overall imbalance on the grid at that instant, as well as on their own imbalances. As a result, it can

be worthwhile for a broker to risk an imbalance if it has the opposite sign from the overall imbalance. This relationship makes the regulating market an interesting strategic element for brokers.

As expected, the figure shows no simple connection between the level of imbalance and the corresponding reward or penalty to the broker. Some general trends, however, can be observed: undersupply (lower half of the figure) generally leads to higher penalties than the rewards afforded by oversupply (upper half of the graph). This is reflective of the asymmetric balancing mechanism employed in Power TAC (de Weerd, Ketter, and Collins 2011), and it illustrates the difficulties in explicitly using the balancing market as part of a strategy. AstonTAC, Mertacor, and to a lesser extent LARGEpower, are mostly well-balanced. They sometimes do incur sizable penalties, however, as a result of small errors on large volumes. MinerTA and SotonPower generally make larger mistakes on smaller volumes, and these mistakes are more favorably distributed between over- and undersupply, leading to smaller overall balancing penalties or even to rewards in the balancing market. UTest makes large, systematic errors on large volumes, leading to hefty penalties.




Figure 5: Balancing market positions. Each point represents one broker/game combination; larger symbols indicate higher overall volumes. Notice the log-log scaling.

To gain a deeper understanding of the dynamics of these games, we finally looked at results in individual games. These analyses shed further light on broker performance differences, and they suggest interesting questions for future research. For the sake of brevity, we discuss only selected examples here. Figure 6 shows the relationship between purchase volume and price in the wholesale market for two complete example games. Interestingly, two very different patterns emerge: the game in the left panel is characterized by the expected upward sloping relationship between volume and price. As higher-cost providers of electricity need to be activated to cover demand, prices rise. Whereas the game in the right panel shows a general upward tendency, it also shows pronounced signs of **economic regime** (Ketter et al. 2012a) transitions (including an “outlier regime” in the upper right-hand corner of the panel), and of start-game effects (lighter dots on the lower left) as brokers adapt to

the market. From a theoretical perspective, it is interesting to consider the properties of the participating strategies that lead to this behavior. Power TAC enables the replay of individual games to support this type of research.




Figure 6: Wholesale prices as a function of volume. Each dot represents the clearing price for one timeslot, lighter dots are earlier timeslots. Left: default broker vs. MinerTA vs. SotonPower; right: default broker vs. LARGEpower vs. SotonPower

While Figure 6 suggests a somewhat homogenous wholesale market, it is worth noting that wholesale bidding in Power TAC is actually highly complex. In addition to deciding the optimal bid price and volume, a broker also needs to factor the optimal **leadtime** into its decision. Recall that for any given timeslot in which electricity is needed, there are 24 separate but interrelated auctions ranging from 1 to 24 hours in advance. Figure 7 shows the impact of leadtime on prices for two selected games. Both panels show the general downward-sloping price trend for increasing leadtimes that economic theory would suggest. However, there are irregularities in these graphs that underline the importance of reasoning about leadtimes and order limits. First, the graphs explain the high prices that Mertacor and UTest are paying in the wholesale market. In the left panel, Mertacor covered all its electricity demand in the 2-hour-ahead auction at prices that are close to double the prices of longer leadtimes. In the process, it drives up prices for AstonTAC and LARGEpower, while SotonPower controls its exposure to this effect through limit orders. An analogous observation holds true for UTest in the right panel. LARGEpower (left panel) placed all its orders as market orders. This leads to relatively good balancing performance, as only forecasting error and no uncertainty in order filling impact its results, but it is also subject to unpredictable fluctuations in wholesale prices that will hamper its risk management efforts.

Conclusions and Future Work

Our energy-hungry civilization must somehow adapt itself to the availability of renewable sources. This will require a combination of new technology and public policy that works with real people and institutions. The Power TAC project intends to contribute by providing a robust research platform for testing market-based approaches to energy sustainability.

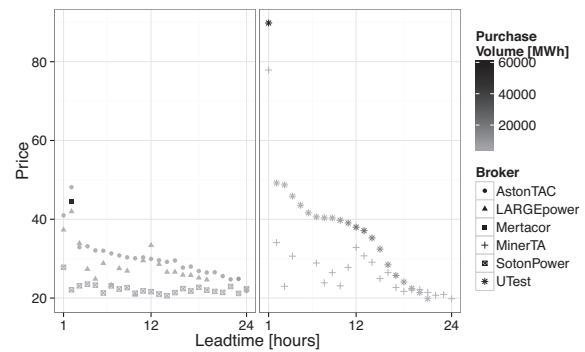


Figure 7: Leadtime effects in the wholesale market; left panel: AstonTAC vs. default broker vs. LARGEpower vs. Mertacor vs. SotonPower; right panel: default broker vs. MinerTA vs. UTest

The annual competition model, along with availability of the platform and a set of working agent implementations, has been an effective driver of research activity in domains as diverse as robot soccer, disaster response, and supply-chain management.

The preliminary analyses we presented are evidence of the realistic macro-level behavior emerging from broker interaction with Power TAC, and of significant performance differences between different approaches to retail electricity trading. Once the participating strategies are fully developed, tools like empirical game theory (Jordan, Kiekintveld, and Wellman 2007) can be leveraged to generate compelling, actionable insights into novel technologies and public policies for future sustainable energy systems.

The Power TAC platform, including the simulator, broker agent framework, log analyzer, and tournament manager, is an open-source project, designed and documented to be accessible to advanced students. Access to the software and documentation, along with a repository of broker agent implementations, will be maintained through the powertac.org website. We look forward to many years of vigorous competition and high-impact research results.

Acknowledgements

Thanks to the Power TAC developers at Minnesota, Aristotle in Thessaloniki, Zagreb, CMU, and Erasmus in Rotterdam, and to all the teams who saw the promise of Power TAC and have been willing to invest in broker development for a rapidly evolving platform.

References

Dash, R.; Jennings, N.; and Parkes, D. 2003. Computational-mechanism design: A call to arms. *Intelligent Systems, IEEE* 18(6):40–47.

de Weerd, M.; Ketter, W.; and Collins, J. 2011. Pricing mechanism for real-time balancing in regional electricity markets. In *Workshop on Trading Agent Design and Analy-*

- sis at the Twenty-second Int'l Joint Conf. on Artificial Intelligence.
- EPRI. 2012. Electric Power Research Institute - Understanding electric utility customers - Summary report.
- Ermon, S.; Xue, Y.; Gomes, C.; and Selman, B. 2012. Learning policies for battery usage optimization in electric vehicles. *Machine Learning and Knowledge Discovery in Databases* 195–210.
- European Commission. 2011. Communication: Energy Roadmap 2050.
- European Commission. 2013. Horizon 2020 Programme.
- Fleten, S.-E., and Pettersen, E. 2005. Constructing bidding curves for a price-taking retailer in the Norwegian electricity market. *IEEE Transactions on Power Systems* 20(2):701 – 708.
- Gottwalt, S.; Ketter, W.; Block, C.; Collins, J.; and Weinhardt, C. 2011. Demand side management - a simulation of household behavior under variable prices. *Energy Policy* 39:8163–8174.
- Jordan, P. R.; Kiekintveld, C.; and Wellman, M. P. 2007. Empirical game-theoretic analysis of the TAC supply chain game. In *Seventh International Conference on Autonomous Agents and Multi-Agent Systems*, 1188–1195.
- Joskow, P. L. 2008. Lessons learned from electricity market liberalization. *The Energy Journal* 29(2):9–42.
- Ketter, W., and Symeonidis, A. 2012. Competitive benchmarking: Lessons learned from the Trading Agent Competition. *AI Magazine* 33(1).
- Ketter, W.; Collins, J.; Gini, M.; Gupta, A.; and Schrater, P. 2012a. Real-time tactical and strategic sales management for intelligent agents guided by economic regimes. *Information Systems Research* 23:1263–1283.
- Ketter, W.; Collins, J.; Reddy, P.; and de Weerd, M. 2012b. The 2012 Power Trading Agent Competition. Technical Report ERS-2012-010-LIS, RSM Erasmus University, Rotterdam, The Netherlands.
- Peters, M.; Ketter, W.; Saar-Tsechansky, M.; and Collins, J. 2013. A reinforcement-learning approach to autonomous decision-making in smart electricity markets. *Machine Learning Journal Special Issue on ECML-PKDD 2012*. Forthcoming.
- Ramchurn, S.; Vytelingum, P.; Rogers, A.; and Jennings, N. 2012. Putting the "smarts" into the smart grid: A grand challenge for artificial intelligence. *Communications of the ACM* 55(4):86–07.
- Reddy, P., and Veloso, M. 2011. Learned Behaviors of Multiple Autonomous Agents in Smart Grid Markets. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI-11)*.
- Reddy, P., and Veloso, M. 2012. Factored Models for Multiscale Decision Making in Smart Grid Customers. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI-12)*.
- Strbac, G. 2008. Demand side management: Benefits and challenges. *Energy Policy* 36(12):4419–4426.
- Sun, J., and Tesfatsion, L. 2007. Dynamic testing of wholesale power market designs: An open-source agent-based framework. *Computational Economics* 30(3):291–327.
- Tesfatsion, L. 2002. Agent-based computational economics: Growing economies from the bottom up. *Artificial Life* 8(1):55–82.
- Voice, T. D.; Vytelingum, P.; Ramchurn, S. D.; Rogers, A.; and Jennings, N. 2011. Decentralised control of micro-storage in the smart grid. In *AAAI Conference on Artificial Intelligence*, 1421–1427.
- Vytelingum, P.; Ramchurn, S.; Voice, T.; Rogers, A.; and Jennings, N. 2010. Trading agents for the smart electricity grid. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*, 897–904. International Foundation for Autonomous Agents and Multiagent Systems.
- Wellman, M. P.; Greenwald, A.; and Stone, P. 2007. *Autonomous Bidding Agents*. MIT Press.