Information Systems Research
Vol. 21, No. 4, December 2010, pp. 688–699
ISSN 1047-7047 | EISSN 1526-5536 | 10 | 2104 | 0688
DOI 10.1287/isre.1100.0316

Research Commentary

Designing Smart Markets

Martin Bichler
Department of Informatics, Technische Universität München, 85748 Garching/Munich, Germany, bichler@in.tum.de
Alok Gupta
Carlson School of Management, University of Minnesota, Minneapolis, Minnesota 55455, alok@umn.edu
Wolfgang Ketter
Department of Decision and Information Sciences, Rotterdam School of Management, Erasmus University,
3000 DR Rotterdam, The Netherlands, wketter@rsm.nl

Electronic markets have been a core topic of information systems (IS) research for last three decades. We focus on a more recent phenomenon: smart markets. This phenomenon is starting to draw considerable interdisciplinary attention from the researchers in computer science, operations research, and economics communities. The objective of this commentary is to identify and outline fruitful research areas where IS researchers can provide valuable contributions. The idea of smart markets revolves around using theoretically supported computational tools to both understand the characteristics of complex trading environments and multiechelon markets and help human decision makers make real-time decisions in these complex environments. We outline the research opportunities for complex trading environments primarily from the perspective of design of computational tools to analyze individual market organization and provide decision support in these complex environments. In addition, we present broad research opportunities that computational platforms can provide, including implications for policy and regulatory research.

Key words: auctions; design; decision support systems; experimentation; smart markets; software agents; platforms; preferences; trading agent competition
History: Vallabh Sambamurthy, Senior Editor. This paper was received on June 30, 2010, and was with the authors 10 days for 1 revision. Published online in *Articles in Advance* November 18, 2010.

1. Introduction

The increase in computational power, advances in user interface design, and the fluidity of electronic communication with the advent and growth of the Internet have realized many of the theoretical conjectures that originated from the seminal article by Malone et al. (1987). The research in computer science, economics, management, and information science is trying to push the envelope on market structure, organization, design, and decision support in increasingly complex and codependent markets in which modern organizations operate. The term *smart markets* has been used in management science (McCabe et al. 1991, Gallien and Wein 2005) to refer to optimization-based markets. We want to generalize this idea and focus on interactive market design, respective decision support, and computational tools to increase the efficiency and robustness of individual smart markets. Our primary focus is on anticipating market needs at various levels so that decision makers such as regulators, manufacturers, and customers are armed with appropriate tools and information to make the best possible decisions. Rapid advances in computational power and evolution of computer networks have enabled the researchers to design complex market structures where *computational intelligence* is desirable to facilitate human decision making. Examples of smart markets or need for synthetic intelligence exist in several markets, such as energy markets, online retailing and negotiations, and FCC spectrum auctions.

The smart markets research that we consider refers to this latter perspective. Smart markets require a wide spectrum of research from optimization techniques to game theoretic formulations, as well as research on individual behavior and preferences to market mechanism and even organizational design. A wide array of methodological developments, from new statistical methodologies to artificial intelligence techniques, also need to take place. In addition, computational research platforms that can model a variety of complex business and economic environments need to be developed. These platforms facilitate exploration of multiechelon systems to study embedded endogenous relationships in a holistic manner. Computational research platforms also *facilitate anticipation* of market needs at a variety of different levels of coarseness so that decision makers such as regulators, sellers, and buyers could be armed with appropriate tools and information to make the best possible decisions, ultimately increasing allocative efficiency of markets and welfare. These approaches have been adopted in many other academic fields (such as social sciences, engineering, and economics) to study phenomena that have been shown to be difficult to analyze analytically. We argue that computational platforms should also be used by information systems (IS) researchers to study electronic market design issues as a complement to pure mathematical models and lab experiments. The purpose of this commentary is to draw attention to a set of particularly fruitful potential areas of research inquiries to which IS researchers can make significant contributions and, therefore, should be involved in setting the direction and research agenda for this stream of research.

For a commentary in an emerging area and for brevity, it is important to define the boundaries of what we are focusing on and why. Smartness, by definition, requires *intelligence*. Therefore, we first define what we mean by intelligence in the context of market design, platform design, and the design of associated decision support tools that form the basic building blocks of smart markets. Note that because our interest is in defining intelligence derived via the use of computing tools, we focus primarily on what could be termed as computational intelligence. However, for expository purposes, in the rest of the document we simply refer to computational intelligence as intelligence. We divide intelligence into two broad categories: (i) instantaneous or real-time intelligence that is used primarily by individual or atomic entities in individual markets and (ii) collective intelligence that refers to market-level intelligence where several atomic entities across markets interact with each other.

The need for *instantaneous* or *real-time intelligence* typically arises in dynamic markets where uncertainty in supply, demand, quantity, quality, etc. creates complex utility functions in multiple dimensions. New multiobject auction designs, such as combinatorial auctions, provide an example for smart market designs where new bidding languages and optimization have led to numerous new research questions in the past few years (Cramton et al. 2006). In such environments, it is often difficult to even express the preferences over the complex decision set adequately, leave alone make a decision. Obviously, increased computational power and improved algorithms for optimal decision making have made significant progress over the decades, and we can solve problems in seconds that might have taken days only a decade ago. However, the markets have evolved as well, and decision makers are now forced into more diverse markets with more complex rules and conditions, making it difficult to use the information that might indeed be available or computable. A prime example of failures of such markets (or exploitation of complex markets due to the difficulty in comprehending its operational rules) was California's electricity market crisis in 2000–2001, with disastrous impact on both the state and companies involved in trading the electricity. In §2, we review the significant strides made in this area and a vast array of open research questions that need to be addressed for creating a higher level of domain centric real-time intelligence.

The *collective intelligence* refers to developing an understanding of complex relationships in any multiechelon ecosystem of codependent and coevolving markets, such as end-to-end supply chains in a competitive market. The enhanced computational powers, open programming paradigms, and advances in multiagent systems research have opened up unprecedented opportunities to study whole ecosystems. The study of these ecosystems provides opportunities to use "neo-Schumpertarian economics" to understand the role of technology and industry structure to explain the evolution of market structures and dynamics of industry change in today's knowledge-based economy (Hanusch and Pyka 2007). Developing open, research-oriented computational platforms has a great promise to facilitate simulation studies at industry level and, as several events in financial markets around the world since 2007 have highlighted, studies to develop appropriate policy and regulatory frameworks. The Trading Agent Competition (http://www.tradingagents.org) is one example of such a computational platform and simulation environment. Computational platforms can also provide an ideal vehicle to study dynamics of competition and strategy at micro level by allowing integration of independent software agents (acting as, for example, firms) to study and validate theoretical results and explore complex economies that cannot yet be analyzed analytically or can be analyzed only with strong and unrealistic assumptions.

In addition, platforms provide ideal environments to one of the thorniest issues in customization and market segmentation—the estimation of consumer preferences—because the ability to test approaches for estimating consumer preferences can be tested against actual preferences that are mathematically or probabilistically defined. The set of research opportunities in this domain are numerous and, with the opportunities to easily develop plug-and-play computational agents, provide an exciting avenue of research opportunity to IS researchers. We discuss these opportunities in §3.

Finally, §4 concludes this commentary with thoughts on the broad set of methodological tools and expertise that can contribute and indeed is desirable, making this research area very exciting.

2. Research Opportunities in Real-Time Intelligence for Smart Markets

As mentioned earlier, the need for real-time intelligence arises in dynamic markets with significant uncertainty about demand and supply. These markets often employ auctions to agree to terms and conditions of a trade. Traditionally, these mechanisms have been studied by using classical game theory and mechanism design, where issues of equilibrium (to ensure efficient outcomes) and incentive compatibility have dominated the theoretical concerns. However, practical smart markets require explicit consideration of computational aspects. This insight has led to many new insights at the intersection of computer science, economic theory, operations research, and information systems. For example, *algorithmic mechanism design* combines ideas such as utility maximization and mechanism design from economics by considering computational constraints to be of central importance for viable economic mechanisms, but at the same time aims for incentive compatibility (Vazirani et al. 2007). IS researchers have considered design enhancements in online auctions by considering actual strategies used by real participants (Bapna et al. 2004) or evaluated different pricing concepts in combinatorial auctions (Bichler et al. 2009). Even in the traditional posted price environments, a seller can use real-time intelligence to dynamically set prices based on consumers' willingness to pay (Bapna et al. 2008). However, for brevity, in the rest of this section, for expository purposes we discuss the research opportunities in traditional dynamic markets (i.e., auction-like environments).

There are numerous research opportunities for IS researchers in the domain of providing real-time intelligence and capability in design. The primary issue from the bidders' perspective is the ability to formulate and express their bidding preferences appropriately. This is a nontrivial task in complex auctions on multiple attributes, multiple units, or items where auctioneers might have incentives to provide limited information and/or the informational environment does not afford the quality of information that bidders can use effectively (Adomavicius et al. 2009). Conversely, the issues for auctioneers revolve around: (i) what information should be revealed, and (ii) what should be the format of a given auction and what should be specific rules of engagement in a given mechanism (e.g., as Bapna et al. 2004 discuss, in online multiunit auctions a variety of rules exists to prioritize the bids at a given bid level). In §§2.1 and 2.2 we discuss each of these research opportunities in more detail.

2.1. Intelligence for Bidders' Decision Making

Broadly, the problem of providing intelligent support for bidders, for a given mechanism, can be divided in two parts: (i) expressing their multidimensional preferences as bids, as opposed to primarily price-based bidding, and (ii) the ability to reasonably accurately predict future state of the auction to understand the quality of their own bids (e.g., how their own bid compares to ending price prediction).

2.1.1. Expressing Complex Preferences. The Internet brought about a large set of innovations and provided bidders the ability to express and exchange information about their preferences in multiple dimensions such as multiple items, multiple units, and possibly multiple attributes of a good or service to be traded in real time. Developing more expressive and yet simple bidding languages is an open research area by using functions or logic formalisms for bidders to describe their preferences in a compact manner. Such bid languages can better match buyers and sellers and consequently create higher allocative efficiency. However, this leads to hard optimization problems, and auctioneers as well as bidders need decision support to harness the new possibilities of such markets. Therefore, such research involves multidisciplinary research teams and solutions to create better bidding languages and new mechanism designs.

Combinatorial auctions are an important type of mechanism that highlights many of the challenges associated with requirements of intelligence for bidder and auctioneer support. Instead of auctioning multiple items in a sequence or in parallel, combinatorial auctions (CAs) allow for bids on bundles of items. Theory and experimental evidence suggest that this can lead to significantly increased efficiency in the case of synergies among items. While many current combinatorial auction implementations are not smart markets, there is a need for intelligence tools in these markets to facilitate better bidding and create higher efficiencies. Many traditional challenges associated with these types of auctions (such as winner determination) have been mitigated with the enhanced computational capabilities and innovative solution techniques to create many real-world applications such as transportation or procurement (Cramton et al. 2006). However, the bidders' problems and need for real-time intelligence still remains a big challenge. These challenges arise for the following reasons:

- A bidder needs to determine their valuations for $2^n - 1$ bundles, where n is the number of items for sale. As an example, this would require elicitation of 1,023 valuations for an auction with only 10 items of interest. This could be thought of as *valuation complexity*.
- Even if the bidders knew their valuations, they would still need to decide how to respond during an auction; for example, when and how much to bid. This can be thought of as *strategic complexity*.

Ideally, research focusing on different auction rules that reduce strategic complexity for bidders is highly desirable along with the need to provide adequate bidding languages and user interfaces to express complex preferences. IS researchers can contribute to both these research streams with their experience in measuring and calibrating electronic systems usage. Bidding languages can exploit parameter space of a trading environment by explicitly allowing the users to represent their preferences; while there have been individual papers on computational aspects of volume discount auctions or multi-attribute auctions (Bichler and Kalagnanam 2005, Goossens et al. 2007), there is little work on optimal configuration of the parameter space and what information should be included or excluded under different market conditions. Much more research is required on knowledge representation for different types of markets and respective smart market designs. Clearly, IS researchers have a lot to contribute in this vibrant area of research.

2.1.2. Understanding Bid Quality. Traditionally, game theory has been used to analytically derive optimal strategies for participants in interactive environments. However, for most auction formats no dominant strategy equilibria exist; therefore, bidders can at best speculate on other bidders' valuations or strategies. Understanding his or her own bid quality is essential for a bidder to place a good quality bid. Judging the quality of a bid becomes a challenging task when the information regarding the mechanism is hard to grasp due to mechanism complexity (e.g., iterative combinatorial auctions where even the revelation of other bidders' bids does not necessarily ease the burden on a given bidder) or information relevant for making a bidding decision is purposefully withheld, as is often the case in procurement auctions where information regarding exact bids might not be provided even in a multi-round iterative auction (e.g., Muni auctions at http://www.grantstreet.com/). IS researchers can develop tools that can support and help bidders in evaluating their bid quality. For example, Adomavicius et al. (2009) provide an interesting computational approach in an opaque auction for bidders to be able to probabilistically estimate a winning bid based on their risk profile. Similar tools can provide bidders a better estimate of valuable information, such as their surplus and winning probability in complex auction and negotiation environments.

IS researchers are already using a variety of tools, such as multi-agent systems and simulations, Monte Carlo and discrete event simulations (e.g., Scheffel et al. 2010), and computational strategy design (Greenwald et al. 2010). However, the area is in its infancy, and significantly more creative research is needed in the area.

2.2. Intelligence for Auctioneers' Design Choices

In addition to combinatorial auctions there have been several contributions on other types of smart markets that allow for volume discount bids or multi-attribute bids. Overall, surprisingly little is known about the design and bidding behavior in these smart markets. A number of fundamental results on equilibrium strategies for selected auction formats, often with strong restrictions on the types of valuations, are negative in the sense that the assumptions for fully efficient mechanisms with a strong game-theoretical solution concept do not typically hold in practical applications. Both research that is based on data from the field and experimental research show that bidders do not necessarily follow optimal strategies as predicted by game-theoretical models (Bapna et al. 2003, Scheffel et al. 2010).

While there is a significant body of theoretical and experimental literature on price-only auctions, the literature on the empirical and behavioral work on complex auctions and other types of smart markets is scarce to say the least. While IS researchers have done some work on decision support tools using alternative feedback mechanisms in combinatorial auctions (Adomavicius and Gupta 2005) or choice of appropriate bid-increment and prediction of future states of the auction in multi-unit auctions (Bapna et al. 2003, 2008), this literature is in its infancy. A significantly higher amount of research needs to be conducted in the area of design choices available to an auctioneer and its impact on behavioral and economic properties of the mechanism using experimental, computational, and simulation-based approaches in addition to theoretical analyses. An example of such research is the experimental work by Scheffel et al. (2010), who show that with linear competitive equilibrium prices a high level of efficiencies can be achieved, even when theoretically no such prediction can be made. IS researchers can also contribute a great amount on participant behavior, antecedents of these behaviors, and moderating and mediating effects of mechanism design choices on these behaviors.

While the research on real-time intelligence for smart markets is exciting and provides significant research opportunities, in the next section we discuss research opportunities using computational platforms that provide an even wider array of promising research opportunities for IS researchers in the dimensions of policy and regulatory research and the holy grail of computational research for the smart market research—the estimation of consumer preferences and its usage in business decision making. Note that significant research opportunities for real-time intelligence exist with computational platforms as well. For example, computational research platforms for Trading Agent Competition for Travel[1] is used for optimizing the product configuration based on partially observable customer demand. Similarly, Trading Agent Competition for Supply Chain Management provides opportunities for research in sourcing and dynamic inventory management decisions. However, the next section focuses primarily on market-level intelligence in the presence of several competitors that can be thought of as atomic agents.

3. Operational and Collective Intelligence

Much of the research on mechanism design has focused on the design of a single isolated event. The analysis of multiple interdependent markets and supply chains turns out to be much more difficult. Neoclassical general equilibrium models have been criticized for having unrealistic assumptions

[1] See http://www.tradingagents.org.

unsuitable for analyzing participant behavior and efficiency of real markets. As mentioned in the introduction, computational platforms and computational experiments can provide significant new opportunities for IS research in exploring a range of interesting and complex problem domains that are both socially and economically important. Computational platforms provide feature-rich and flexible research environments that are low risk, holistic laboratory platforms for developing an understanding of the emergent phenomena arising from complex interactions of market mechanisms with adaptive, self-interested agents, under any technical and economic constraints that future smart markets may impose. These platforms can help us understand the interplay of smart markets and better quantify the symbiotic relationships among various market participants in the dynamic multiechelon economy.

These platforms and tools become possible through the emergence of new, previously unavailable large-scale computational power such as cloud and grid computing. Applied computer scientists have already designed some platforms that have become the focus of research, with *competitions* being organized to bring the research community together and to test the efficacy of both computational and strategic research ideas in a given domain (www.tradingagents.org). Researchers have observed that these platforms and associated competitions have been an effective way to spur innovation (Kearns and Ortiz 2003). On the macro level, these platforms are able to provide clear guidance for policymakers on the capabilities and limitations of open market structures for management of future smart markets in a given domain. On the other hand, individual agents (firms) can use these to learn, create, and test short- and long-term strategies using available information from the overall business ecosystem. The participants in these competitions need to design techniques to manage risk and adapt to changing conditions while concurrently trading in multiple markets. It is hard to do experiments with real organizations, but these experimental platforms provide a powerful vehicle for evaluation and validation of new and creative ideas. Organized competitions such as trading agent competition for supply chain management (TAC SCM) (Collins et al. 2010), and Power TAC[2] (Ketter et al. 2010), along with many related tools, are driving research into a range of interesting and complex domains that are both socially and economically important. Because such experimental markets allow market structures to be evaluated under a variety of real-world conditions under competitive pressures, these platforms can also be used to effectively uncover potential hazards of proposed market designs in the face of strategic behaviors on the part of the participating agents, helping policy makers in policy and regulation design. For instance, there were multiple opportunities for agents to manipulate the game in unintended ways (Ketter et al. 2004). Similar phenomena can be observed in real-world supply chains, where certain policies can be an advantage to some suppliers and lead to unintended long-term effects on the market.

Multiechelon markets, such as supply chains, should not be viewed as monolithic entities that can be centrally optimized, but instead consist of multiple self-interested entities each operating according to its own objectives and policies. For example, while each real-world supply chain exhibits its own peculiarities, we need platforms that capture major sources of complexity common to many supply chains, shielding researchers from less relevant idiosyncrasies. IS researchers should become central players in designing these platforms and the entities around these platforms. Smart market entities could be modeled as autonomous agents that concurrently compete with each other in environments that are subject to both exogenous and endogenous sources of uncertainty. These platforms reflect many desirable characteristics of the real world because agents act autonomously to maximize their expected utilities in an environment that is highly dynamic, with only partially observable parameters, and strongly affected by competitive forces.

3.1. Management and Control of Platform Mechanisms

There are numerous research questions regarding operational and collective intelligence that can be answered with the help of these computational platforms; however, due to space limitations we focus

[2] See http://www.powertac.org.

only on the important questions of policy and regulatory recommendations, as well as on the coordination among different markets.

3.1.1. Policy and Regulatory Recommendations. Platform-based research will allow exploration of policy and regulatory recommendations for existing and emerging markets—an area that has lacked any serious research efforts beyond rhetorical debates. Researchers can generate an improved understanding of technology-based markets and interactions to inform policymakers regarding the capabilities and limitations of open market structures because the problems in such markets are about to become even more complex. As the devastating and unintended consequences of the California energy crisis of 2000 and 2001 highlighted, the agents in any trading environment will try to exploit mechanisms weaknesses in an unintended manner. Theoretical analysis, due to its mathematical constraints and assumptions, is unlikely to discover design problems or limitations. Computational research platforms can explore and provide warnings regarding potential loopholes due to limitations of design, implementation, or regulatory policies.

For instance, electricity production and distribution systems are already complex adaptive systems that need to be managed in real time to balance the load of an electricity grid. Moreover, the electrical energy market is about to undergo a significant shift in usage pattern, where a large amount of potential individual electrical energy consumption coming from electric vehicles as opposed to just homes and/or offices. Furthermore, many markets such as electricity markets are currently undergoing a transition from centrally regulated systems to decentralized but interdependent markets. These transitions are very risky because we do not have sufficient experience in setting up decentralized energy systems and predicting their effect on the economy. As mentioned earlier, lack of appropriate policies can cause major damage while structurally changing the design and interactions in any complex real-world market. The California energy market debacle and the related collapse of Enron have demonstrated that the success of competitive electricity markets crucially depends on market design, demand response, capacity reserves, financial risk management, and reliability control along the electricity supply chain. Similar arguments could be made about a variety of markets that involve multiple markets at multiple levels with competing interests that cannot all be simultaneously analyzed in any mathematically tractable way.

Another interesting factor is in the burgeoning era of crowdsourcing, or the newly coined term *shared-sourcing*. In many emerging smart markets, such as energy, the line between the traditional roles of producers and consumers is already starting to blur, and so-called *prosumers* are emerging. In other words, the consumers of energy can become producers of energy as new technologies of storing and producing energy become mainstream. Similar advances in other areas (e.g., consumer-to-consumer markets) can change the roles of an entity from transaction to transaction, and the producers in one instance can become consumers in another, and vice versa, even for the same underlying commodity.

IS researchers have a great opportunity to study how such potential shifts in the environment affects the organization and the participants in these markets. Our belief is that IS researchers, besides artificial intelligence and economics, should be part of the community that designs these new competitive simulation platforms that implement market-based management structures to study future smart markets. These platforms will challenge research teams to create agents, or possibly agent-assisted decision support systems for human operators, which can operate effectively and profitably in direct competition with each other. Research teams with a variety of specific interests thereby could contribute to the development of reliable and efficient automation technologies for efficient trading on the retail and wholesale level. At the same time, they are challenged to exploit the structure of the smart market and to design environments (by iteratively adjusting them) to defeat counterproductive strategic behaviors.

In summary, platforms will allow researchers to study policy issues in markets at different levels through the development of competitive, agent-based automation strategies. These platforms allow complex market structures to be evaluated in a risk-free environment under a variety of real-world conditions ranging from normal to extreme. The competitive design can effectively uncover potential hazards of proposed smart market designs in the face of strategic behaviors on the part of the participating agents; countervailing policies and regulations can thus be designed more effectively. We next identify and discuss one of the most intriguing areas of research that involves studying the dynamics of interaction among various mechanisms that form the basis of complexity in a given market domain.

3.1.2. Coordination Among Mechanisms. The rapid rise of Internet-enabled business interactions makes the supply chain management domain, like many other real-world multiechelon problem areas, increasingly challenging for human decision making. Besides traditional challenges, such as inventory management and shop-floor optimization, most of the supply chain issues currently deal with the coordination of supply chain operations, including coordination with suppliers, various functional areas within the firm, distributors, etc. Because the incentives of various entities within a firm's supply chain could have conflicting goals, coordination usually requires development of mechanisms that can help align the incentives appropriately. Given that the evaluation of these new approaches to decision making is very difficult in isolation, researchers need to have flexible experimental platforms to study and evaluate market participants' relationships in end-to-end supply chains with various methodologies. Therefore, the emerging computational platforms should place a premium on effective coordination of decisions affecting multiple markets and internal resources. For instance, in supply chain management, inventory planning is complicated by the fact that a given part might be used in multiple products, and a shortage of a particular part can prevent an agent from participating in significant segments of the customer market. IS researchers have much to contribute to this difficult coordination problem because they can design, implement, and test different real-world coordination strategies (Collins et al. 2009, 2010).

3.2. Platform-Based Multiechelon Preferences

The human mind has limited cognitive capacity, so humans tend to make decisions using rules of thumb, or heuristics, which stem from their own experiences (Simon 1979). An important area of study is the design and implementation of (artificially) intelligent agents and decision support tools that can effectively assist humans with their decision making efforts (Wooldridge and Jennings 1995), particularly in information-rich and time-critical environments. As an example, consider the Dutch Flower Auctions (DFA), where human decision makers have to bid in an environment with up to 40 auctions simultaneously, where each auction lasts for 3–5 seconds (Kambil and van Heck 1998). Similarly, the amount of data that must be processed by bidders in widespread spectrum auction designs in a relatively short time requires decision support. Furthermore, in multiechelon systems the preferences of the same individual, e.g., price elasticity, might be different in different decision or environmental regimes—for example, in situations of oversupply vs. scarcity (Ketter et al. 2009). Software agents that can mimic human behavior (or perhaps more efficiently adapt) have the potential to improve bidder performance significantly, because they open possibilities for automating, augmenting, and coordinating decision processes. These agents can act on behalf of a user with some degree of independence or autonomy, employing some representation of the user's goals and preferences. Researchers should focus on enhancing the adaptive learning component of such agents. These agents could predict the appropriate next steps, helping to speed up and improve the quality of a user's overall decision process. Therefore, research is needed to develop and evaluate highly personalized software agents that complement the cognitive and computational capacity of humans, while leveraging the experience and contextual knowledge of seasoned decision makers. These agents will collaborate with their users to gather and present information and recommend action.

To work effectively and efficiently on all levels of decision making in multiechelon market environments, these agents must learn the preferences of their users with respect to the decision context. Preferences and their influence on decisions have been studied extensively in economics and related fields, but many important problems remain untouched. One of the big challenges is how to adapt the existing static representation methods to dynamically model user preferences in a compact form. This is important for fast-evolving domains, such as smart energy markets or Internet-enabled businesses.

The challenge is that current market design approaches assume full knowledge of a bidder's preferences, but *preference elicitation* is a formidable task for several reasons. First, in cooperative environments, users face well-documented difficulties of accurately articulating preferences (Shiv and Fedorikhin 1999). For example, in combinatorial auctions, bidders would need to know their valuation for an exponential number of bundles. Experiments have shown that bidders without tool support evaluate and bid on only a small amount of bundles (Scheffel et al. 2010). This is a significant source of inefficiency, and much higher efficiency could be achieved through appropriate decision support. Second, even when focusing the users' attention using well-established preference elicitation methods, such methods tend to be cumbersome and give inconsistent results. Finally, in competitive environments, users might not want to divulge their preferences. Mechanism design theory deals with such questions, but little is known about incentive compatibility in repeated and interrelated markets.

Figure 1 presents an agenda that offers numerous research opportunities for IS researchers in designing intelligent software agents for decision making in smart markets, and they avoid many of the shortcomings from previous research. We partition the research agenda into three overlapping phases and suggest novel work in analytical, empirical, and experimental research in agent-based multiechelon smart market platforms. To act rationally, the agents must choose among different actions and means of expression; to intelligently support the actions of decision makers, they must understand and respond to the users' choices. Therefore, the agenda includes historical data analysis and modeling (Phase 1), autonomous decision automation including dynamic market and decision modeling (Phase 2), and agent-assisted real-time decision support (Phase 3).

Phase 1: Preference model representation and elicitation. Business decisions represent choices that combine a variety of observable factors we call *market dynamics* with explicit goals and implicit preferences. These decisions are reflected in transactions. For example, supply and demand in energy markets are characterized by a range of exogenous influences, including daily, weekly, and seasonal demand and supply variations, as well as weather, fuel price, political factors, and competitor behaviors. Phase 1 research activities should be concerned with learning implicit preferences by observing market dynamics and actual transactions and by constructing a model that can account for these observations (see Figure 1). Because currently the data, with high volume and frequency, are significantly ahead of theory, and a rational agent is trying to process this large amount of data to maximize the expected utilities, IS researchers must develop novel statistical methods

Figure 1 **Dynamic Preference Modeling and Decision Recommendation**

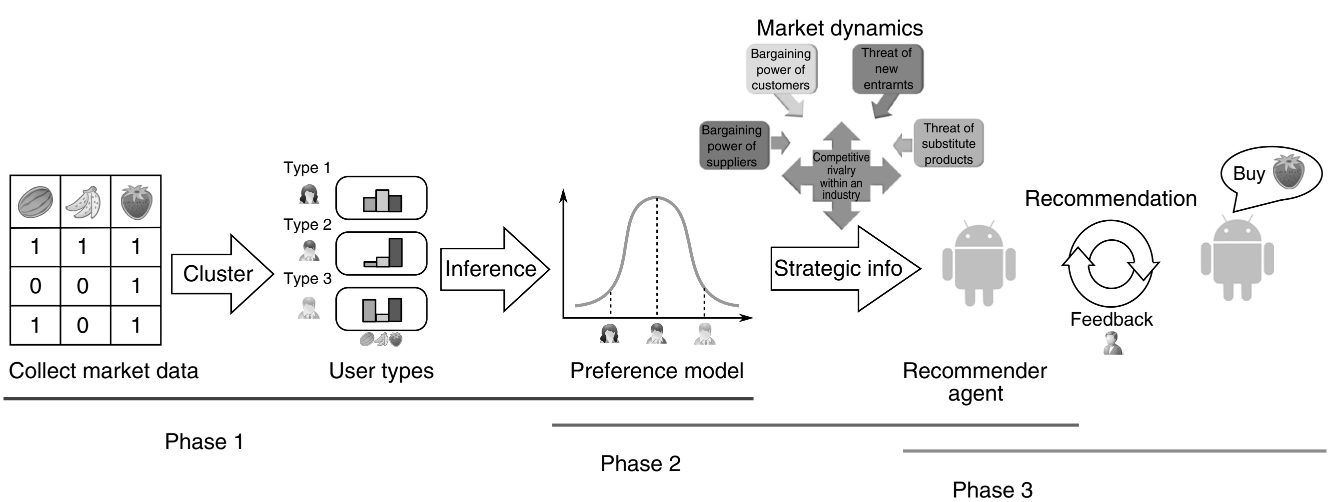

and probabilistic machine learning algorithms that can create insight into existing and future smart markets. This offers IS researchers exciting opportunities to develop new theoretical frameworks, such as how to learn new models of latent behavioral characteristics of the different market participants groups offline, and perhaps use the results in Phase 2 to dynamically predict future market regimes and own decisions in real time. An understanding of customer behavior through such latent behavioral models can lead to significant changes in issues such as product or service innovation and design, dynamic pricing strategy, and distribution strategy.

Phase 2: Autonomous model update and decision support. Many business decisions are made based not on current conditions but on anticipated future conditions. Therefore, an effective decision support agent must be able to identify current and anticipate future market conditions and update the preferences of the decision models appropriately (Ketter et al. 2009). IS researchers must acquire advanced skills in probabilistic modeling so that they are able to develop new machine learning methods, which in turn are able to process the huge amounts of data available and turn it into highly useable information that can be used by automated decision models. Results from scientific computational platforms, such as the various Trading Agent Competitions (TAC), have been successfully translated in the online advertising industry. Another example is Microsoft's experimental platform (http://exp-platform.com/default.aspx), which accelerates innovation through trustworthy experimentation with live users with great success. The platform enables testing new ideas quickly using the best-known scientific method for establishing causality between a feature and its effects: randomized experimental design. The basic methodology in controlled experiments is to expose a percentage of users to a new treatment, measure the effect on metrics of interest, and run statistical tests to determine whether the differences are statistically significant, thus establishing causality. There is a great opportunity for many more smart research platforms like this, because firms must understand the aggregate evolution of customer preferences and adjust their decision models and Web site interfaces accordingly to increase customer satisfaction and other KPIs.

Phase 3: Real-time, agent-assisted decision support. Two problems have and continue to hamper the adoption of agents in industry (Maes 1994). The first is *competence*: to be helpful, an agent must acquire relevant knowledge and must learn from the user in what circumstances to make recommendations. The second is *trust*: the agent must behave in a way that inspires trust in its user. A successful agent-based approach must provide well-grounded suggestions and leave the human decision maker in final control. IS researchers should contribute to this research by designing and developing adaptive learning agents to assist human decision makers in real-time, multi-echelon market environments.

If an agent is learning user preferences, it must not only communicate these predictions to the user but also must know the extent to which the user disagrees with its predictions. To address this question, computer science researchers typically have used ad hoc combinations of various parametric and nonparametric prediction modeling methods to create the *best fit* for a given context. However, establishing a communication paradigm where an agent understands how well its recommendations were received and how to improve its own performance for a given user is an open question. For example, should an agent provide an "optimal" answer given a set of predefined parameters, or should it present a list of alternatives that a human decision maker could choose from? Furthermore, the decisions regarding timeliness or desirability of a recommendation (i.e., whether or not the recommendation would be perceived as nonintrusive) have not been studied much and are significant research areas where IS researchers can make huge contributions, given the IS literature in technology acceptance and usage. Finally, *adjustable autonomy* (i.e., the degree of autonomy that an agent or automated decision support system should have) is a completely open research question, especially in dynamic and complex market environments where specifications of a priori optimization constraints are challenging due to the necessity of subjective decision making.

One of the key virtues of such an agent preference framework is that these models can be refined via multiple kinds of *feedback*. Because the suggested models are typically probabilistic and capture the relationship between high-level preferences and transaction decisions, the preference models can be used to predict a customized transaction decisions for each user. Discrepancies between these predictions and actual choices of users provide a natural error signal for probabilistic models, which can be used to update model parameters in real time. The exact methods and approaches to provide these feedbacks and the computation of errors is in itself challenging and is an ongoing research area that needs significant attention.

In summary, to make decision in a desired way, agents not only must understand the consequences of their actions but also need a policy for choosing what to recommend to the decision maker. That policy could consider instantaneous or long-term effects of choices, but it must have a means of evaluating and computing these effects. Preferences achieve this and are the key for agents to make decisions in a rational way. Therefore, agents must be able to capture temporal components of decision makers' preferences over time in a nondisruptive, indirect manner and update them with explicit feedback so that the user can gain confidence and trust in an agent's abilities. For instance, in multi-unit sequential Dutch auctions, an agent could give iterative recommendations for an auctioneer to set the starting price and minimum purchase quantity of each auction, taking various endogenous (such as bidder population and historical data) and exogenous (such as news and weather) information into account.

4. Conclusions

While there have been a number of contributions on algorithmic mechanism design and optimization in smart markets in recent years, the experimental and decision support literature is still in its infancy. Beyond basic (im)possibility results, the goal is to achieve satisficing, rather than optimal or fully efficient solutions, by considering cognitive and informational aspects of market participants.

Overall, we believe that there exists a significant research opportunity to design smart market mechanisms for complex multi-item, multi-unit, and/or multi-attribute markets using computational tools and in designing decision support tools for markets that impose cognitive challenges for humans. While research on smart markets will continue to be interdisciplinary, the traditional IS research areas such as the design of decision support systems and the modeling of user behavior can play an important role in solving relevant problems in this domain. The area provides many opportunities where IS researchers can contribute in the development of real-world markets with a significant impact on the market participants and the economy.

While theory provides central foundations and boundaries of what is possible, the design of smart markets requires engineering, decision support, and experimental work in order to *put theory to work*. This is very similar to the way relational algebra provides a formal foundation for relational databases, but data engineering comprises many more techniques to provide robust and efficient database technology. Our argument is that while trying to model complex environments, theory should be used at a much more granular level to provide guidance for architecture of platforms rather than at a finer level of granularity such as individual interactions. For example, Ketter et al. (2009) use the idea of economic regimes to capture the current and anticipated balance between supply and demand to define broad guidance for strategic decisions by an atomic entity. The individual interactions then emerge and provide the atomic entity with real-time information to make its tactical decisions. Furthermore, the research in smart markets should always take both technological and knowledge maturity into account for long-term evolution of intelligence in a given market. One danger that computational platform research can potentially suffer from is the inherent desire of researchers to reap instant rewards (such as winning in a competition) as compared to long-term principled approaches that might provide more sustainable insights. In addition, behavioral models and empirical analyses are keys to understanding and evaluating new market designs. Computational, field, and lab experiments will play a pivotal role along with formal methods to design efficient markets, which are robust against various types of speculation and manipulation that we have experienced it in the past.

References

Adomavicius, G., A. Gupta. 2005. Toward comprehensive real-time bidder support in iterative combinatorial auctions. *Inform. Systems Res.* **16**(2) 169–185.

Adomavicius, G., A. Gupta, D. Zhdanov. 2009. Designing intelligent software agents for auctions with limited information feedback. *Inform. Systems Res.* **20**(4) 507–526.

Bapna, R., P. Goes, A. Gupta. 2003. Replicating online yankee auctions to analyze auctioneers' and bidders' strategies. *Inform. Systems Res.* **14**(3) 244–268.

Bapna, R., P. Goes, A. Gupta, Y. Jin. 2004. User heterogeneity and its impact on electronic auction market design: An empirical exploration. *Management Inform. Systems Quart.* **28**(1) 21–43.

Bapna, R., P. Goes, A. Gupta, G. Karuga. 2008. Predicting bidders' willingness to pay in online multi-unit ascending auctions: Analytical and empirical insights. *INFORMS J. Comput.* **20**(3) 345–355.

Bichler, M., J. Kalagnanam. 2005. Configurable offers and winner determination in multi-attribute auctions. *Eur. J. Oper. Res.* **160**(2) 380–394.

Bichler, M., P. Shabalin, A. Pikovsky. 2009. A computational analysis of linear-price iterative combinatorial auctions. *Inform. Systems Res.* **20**(1) 33–59.

Collins, J., W. Ketter, M. Gini. 2009. Flexible decision control in an autonomous trading agent. *Electronic Commerce Res. Appl.* **8**(2) 91–105.

Collins, J., W. Ketter, M. Gini. 2010. Flexible decision support in dynamic interorganizational networks. *Eur. J. Inform. Systems* **19**(4) 436–448.

Collins, J., W. Ketter, N. Sadeh. 2010. Pushing the limits of rational agents: The trading agent competition for supply chain management. *AI Magazine* **31**(2) 63–80.

Cramton, P., Y. Shoham, R. Steinberg. 2006. Introduction to combinatorial auctions. P. Cramton, Y. Shoham, R. Steinberg, eds. *Combinatorial Auctions*. MIT Press, Cambridge, MA.

Gallien J., L. Wein. 2005. A smart market for industrial procurement with capacity contraints. *Management Sci.* **51**(1) 76–91.

Goossens, D. R., A. J. T. Maas, F. Spieksma, J. J. van de Klundert. 2007. Exact algorithms for procurement problems under a total quantity discount structure. *Eur. J. Oper. Res.* **178**(2) 603–626.

Greenwald, A., K. Kannan, R. Krishnan. 2010. On evaluating information revelation policies in e-marketplaces: A Markov decision approach. *Inform. Systems Res.* **21**(1) 15–36.

Hanusch, H., A. Pyka. 2007. Principles of neo-Schumpeterian economics. *Cambridge J. Econom.* **31**(2) 275–289.

Kambil, A., van Heck, E. 1998. Reengineering the Dutch Flower Auctions: A framework for analyzing exchange organizations. *Inform. System Res.* **9**(1) 1–19.

Kearns, M., L. Ortiz. 2003. The Penn-Lehman automated trading project. *IEEE Intelligent Systems*. IEEE Computer Society, Washington, DC, 22–31.

Ketter, W., J. Collins, C. A. Block. 2010. Smart grid economics: Policy guidance through competitive simulation. Technical Report ERS-2010-043-LIS, RSM Erasmus University, Rotterdam, The Netherlands.

Ketter, W., J. Collins, M. Gini, A. Gupta, P. Schrater. 2009. Detecting and forecasting economic regimes in multi-agent automated exchanges. *Decision Support Systems* **47**(4) 307–318.

Ketter, W., E. Kryzhnyaya, S. Damer, C. McMillen, A. Agovic, J. Collins, M. Gini. 2004. Analysis and design of supply-driven strategies in TAC-SCM. Trading Agent Design Anal. Workshop, Columbia University, New York, 44–51.

Maes, P. 1994. Agents that reduce work and information overload. *Comm. ACM* **37**(7) 31–40.

Malone, T. W., J. Yates, R. I. Benjamin. 1987. Electronic markets and electronic hierarchies. *Comm. ACM* **30**(6) 484–497.

McCabe, K., S. Rassenti, V. Smith. 1991. Smart computer-assisted markets. *Science* **254**(5031) 534–538.

Scheffel, T., A. Pikovsky, M. Bichler, K. Guler. 2010. An experimental comparison of linear and nonlinear price combinatorial auctions. *Inform. Systems Res.*, ePub ahead of print February 1, http://isr.journal.informs.org/cgi/content/abstract/isre.1090.0267v1.

Shiv, B., A. Fedorikhin. 1999. Heart and mind in conflict: The interplay of affect and cognition in consumer decision making. *J. Consumer Res.: An Interdisciplinary Quart.* **26**(3) 278–292.

Simon, H. A. 1979. Rational decision making in business organizations. *Amer. Econom. Rev.* **69**(4) 493–513.

Vazirani, V. V., N. Nisan, T. Roughgarden, E. Tardos. 2007. *Algorithmic Game Theory*. Cambridge University Press, New York.

Wellman, M. P., A. Greenwald, P. Stone. 2007. *Autonomous Bidding Agents*: *Strategies and Lessons from the Trading Agent Competition* (*Intelligent Robotics and Autonomous Agents*). MIT Press, Cambridge, MA.

Wooldridge, M., N. R. Jennings. 1995. Intelligent agents: Theory and practice. *Knowledge Engrg. Rev.* **10**(2) 115–152.