Research Statement

David C. Parkes

September 2007

My research bridges between *computer science* and *microeconomic theory*, specifically around the intersections among *multi-agent systems*, *game theory* and *optimization*. The area of multi-agent systems (MAS), a sub-field of artificial intelligence (AI), studies the computational aspects of distributed, rational decision making. Game theory studies the behavior of self-interested agents such as individuals or firms in situations of strategic interdependence. Optimization, as studied for instance in operations research, seeks to develop computational methodologies for, and characterizations of, solutions to problems of maximizing an objective subject to constraints. In working across these general areas, my work has a common goal, which is

> *... to enable optimal system-wide decisions in distributed, possibly dynamic environments populated by multiple participants, each with local information, local capabilities, and individual objectives.*

In bringing together the insights and techniques and handling the cross-cutting constraints of computer science and microeconomic theory, I have focused in particular on the following three topics: (a) *computational mechanism design*, which addresses both computational and incentive constraints and solves problems of optimal coordination in environments with distributed information and self-interest; (b) *efficient combinatorial markets*, which address problems of multi-agent preference elicitation and combinatorial optimization, and find application to electronic commerce and within supply chains; (c) *incentive-compatible social computing*, which seeks to leverage economic methods *within* computational systems and embraces the many distributed systems that are both computational and economic in their characteristics, including peer-to-peer, open science grids, and collaborative Web applications. I will touch on each of these topics in turn.

Computational Mechanism Design. Mechanism design (MD), the origins of which can be traced back to the late '60s and '70s, in a time of great debate about market economies vs. command-and-control economies, offers principles for the design of provably optimal coordination mechanisms in the presence of private information and self-interest. The dominant paradigm is that of incentive compatible direct-revelation mechanisms: each agent first reports a claim about its private information, such as its utility for different outcomes, and then some "center" solves an optimization problem and chooses an outcome and payments. A mechanism is incentive compatible when truthful reports of private information form a game-theoretic equilibrium.

Computational mechanism design (CMD) extends MD to include considerations of computational complexity and communication complexity, and addresses the problem of

> *... making optimal system-wide decisions in multi-agent systems, despite the individual self-interest and local information of agents, and subject to limited computational resources.*

In pursuing this agenda, I have worked to expand the reach of MD to rich and complex domains. An important aspect of this relates to the design of incentive mechanisms for *dynamic environments* in which agents may arrive and depart and have stochastic local state. Some highlights include:

- I introduce an online mechanism for MAS with dynamic arrivals and departures, formalizing the efficient social choice problem in this setting as a Markov decision process. This collaboration with Singh [PS03, PSY04] provides, for example, an auction to allocate an uncertain supply of last-minute theater tickets to a stream of variably-impatient tourists. In the process, I identify a fundamental tension in achieving dominant-strategy incentive compatibility in dynamic environments and settle for a weaker solution concept.

- I introduce a unified framework for CMD in dynamic environments that exhibit both dynamic arrival-departures *and* agents with stochastic private state. This work is in collaboration with graduate student Cavallo and collaborator Singh [CPS06, CPS07] and also leverages the beautiful work of Bergemann and Välimäki [BV07] on dynamic Vickrey mechanisms. The incentive mechanisms that we develop can find application to many problems, ranging from Web monetization, to dynamic supply chains, to the coordination of loosely coupled computational processes.

- I characterize policies that can be implemented in a truth-revealing *dominant strategy* equilibrium in a restricted domain. Developed in collaboration with Hajiaghayi, Kleinberg and Mahdian [HKP04, HKMP05], this characterization requires *decision monotonicity* of a policy in response to unilateral changes in the report of a single agent and extends the classic theory of truthful MD by leveraging constraints on the temporal manipulations available to agents. In recent work with undergraduate student Duong, I demonstrate that optimal policies need not be monotonic and provide a practical, sample-based stochastic optimization framework, in which an *approximate* decision policy is transformed into a monotonic (and thus truthful) policy [PD07]. This is achieved through a new paradigm of "output ironing" in which decisions for which monotonicity cannot be established are canceled.

Another issue in expanding the reach of MD to rich and complex domains, also addressed in my work, is the fundamental mismatch between the centralization of computation implicit in MD theory and the distributed computational models of MAS and networked systems. With graduate student Shneidman [SP04, PS04b], I define the problem of *faithful distributed implementation*. This expands on Feigenbaum's agenda [FS02, FSS06] of distributed algorithmic mechanism design, by requiring that an entire algorithm – message passing, computation, and information revelation – be brought into an *ex post* Nash equilibrium. In recent work, Feigenbaum and colleagues [FRS06] also adopt this faithfulness requirement, and use our proof techniques to bring a richer class of routing policies into equilibrium.

In a related collaboration with Petcu and Faltings [PFP06], I provide the first synthesis of a *distributed constraint optimization algorithm* (DCOP) with the methods of faithful distributed implementation. This brings the methods of MD to bear on DCOPs, which have until now been restricted in application to *cooperative* multi-agent systems, for instance those controlled by a single firm, or instantiated by a single protocol designer.

Efficient Combinatorial Markets. Combinatorial auctions (CAs) address allocation problems in which agents have values on *sets* of goods; e.g., *complements*, "I want A and B together" and *substitutes*, "I want A or B." CAs find numerous applications, including to MBA course registration, the sourcing of goods and logistics in the supply chain, and to the allocation of public goods such as wireless spectrum [CSS06]. I study *combinatorial markets*, in which

> ... *market participants can precisely express, directly or indirectly, information about their private valuations, and the market clearing problem is one of combinatorial optimization.*

In addition to CAs, combinatorial markets include *combinatorial exchanges*, in which agents are engaged in multi-way trade including swaps and other forms of contingent trade, and *multi-attribute auctions* in which the market configures attributes of goods (delivery terms, quality, alternate specifications) in addition to determining the pattern of trade.

One of my most important contributions in this area is to provide a *primal-dual theory for incentive-compatible, ascending-price CAs.* This agenda was initiated in my Ph.D. thesis [Par01]. I argued there that *preference elicitation* is the most important bottleneck to the deployment of CAs and called for *indirect* mechanisms in which agents refine bids across multiple rounds. With my Ph.D. advisor, Lyle Ungar, I introduced a primal-dual auction algorithm (*i*Bundle), that was the first efficient, ascending price CA for a straightforward bidding strategy. More recently, it has been demonstrated that *i*Bundle, when interpreted in sealed-bid form, enjoys the Vickrey auction's robustness to manipulation for a subclass of combinatorial auction problems, while avoiding problems associated with revenue and collusion on general instances [AM02]. A proxied variant on *i*Bundle has been advocated by others for use in the FCC's wireless spectrum auction [ACM06].

In collaboration with Mishra [MP07], I was able to adopt a duality-based characterization of the information required to achieve the Vickrey outcome, developed jointly with graduate students Lahaie and Constantin [LCP05], to extend *i*Bundle and provide the only known *ascending-price, generalized Vickrey auction* for CAs. This auction, which brings straightforward bidding into an *ex post* Nash equilibrium, has been studied by economists in experimental labs [CT05] and is also the subject of agent-based experiments [SSB07]. It addresses the long-standing agenda in economics of bringing simple demand-revealing processes into an equilibrium within an efficient market, in this case for combinatorial domains.

In related work, I also provide a general methodology to adapt algorithms in *computational learning theory* to provide efficient elicitation protocols for CAs. Whereas Nisan and Segal [NS06] show that any efficient CA requires worst-case communication that scales exponentially in the number of items and the number of agents, my results with graduate student Lahaie [LP04] provide *polynomial query complexity* when one also allows the number of queries to depend on the minimal representation size in some bidding language. This neatly turns the tables, making the relevant question one of identifying bidding languages that are concise for problems of interest, and suggests a practical architecture for electronic markets.

Many resource allocation domains (e.g. wireless spectrum, airport takeoff and landing slots, open science grids) have private information on both sides of the market– with multiple buyers and multiple sellers –and therefore require combinatorial exchanges rather than CAs. Initiated as a semester-long project in my graduate seminar, I have designed and validated the *first iterative, fully expressive* combinatorial exchange (ICE). This research was spearheaded in recent years by graduate students Lubin and Juda [PCE+05, LJC+07]. ICE provides a bidding language with tree-based semantics that is concise in domains of interest, allows efficiency guarantees despite employing only linear prices, and scales to economically interesting problems.

Incentive-Compatible Social Computing. By *social computing*, and modifying a definition due to IBM's social computing research group, I refer to

> ... *systems that support the gathering, representation, processing and dissemination of information, the information and (part of) the system itself being distributed across, and controlled by, multiple individuals, teams, and organizations.*

Examples of systems that fall within this definition include recommender systems, collaborative systems such as Wikipedia that rely on information gathering and processing by individuals, peer-to-peer systems, peer production, open science grids, and shared research testbeds. My interest is in welfare-optimizing, *incentive-compatible* social computing, which places an emphasis on methods to schedule constrained resources and provide incentives to promote the efficient contribution of resources, information and effort by participants. Viewed in this way, many of the challenges in social computing look like economic problems: How can incentives be provided to promote the contribution of upload bandwidth in return for downloading files? Why should one organization share its compute resources with another on a computational grid?

With graduate student Ng and collaborators, I experimented with deployments of artificial markets on a sensor network testbed at Intel Research Berkeley [CBA$^+$05] (in continual operation since Spring '05) and on the PlanetLab overlay network [ACN$^+$07]. The sensor network experiment was especially informative because it generated some real-world examples of users manipulating a deployed mechanism for individual benefit [NBC$^+$05]. Building on these lessons, which are summarized in a paper written with graduate students Shneidman, Ng and collaborators [SNP$^+$05], I am now involved in the *Egg* project (Environment Computing for Global Grids) [BHH$^+$06]. This is a multi-year collaboration with high-energy particle physicists (Saul Youssef[1] at Boston University and John Huth at Harvard), colleague Prof. Margo Seltzer, and graduate students Kang and Ng. In Egg we seek to provide efficient resource allocation while recognizing the importance of ease of use and the need to support organizational policy. Novel design features include the use of multiple currencies and macroeconomic controls to support virtual organizations, and an open, extensible architecture, providing strategyproofness properties without mandating a particular auction protocol [KP07]. What is exciting from my perspective is that Egg, if successfully deployed, will provide an extremely rich platform for all kinds of research in incentive-based computing.

In the context of peer-to-peer systems and in collaboration with graduate student Shneidman and Massoulie [SPM04], I disprove faithfulness of the BitTorrent protocol for peer-to-peer file sharing by finding a vulnerability using a backtracing technique. A new collaboration with graduate student Seuken, Johan Pouwelse and the *Tribler* peer-to-peer project, offers an opportunity both to understand the behavior of users in peer-to-peer communities– we recently ran a large-scale experimental study –and also to develop more efficient peer-to-peer economies through careful design.

Closing Comments. I think about my overarching goal as a researcher as one of

> ... *understanding how to combine markets and optimization in addressing problems of agency and in providing efficient behavior in dynamic, complex environments.*

[1]Part of the seed idea for Egg is *Pacman*, designed by Youssef and the *de facto* standard for software installation on U.S. science grids (2.7 million+ Pacman downloads to 50+ countries, with 4.5 million+ software installations created in '06.)

I am motivated in this effort by the awesome power of markets in coordinating human activity, by the scale of problems solvable by computational methods today and the contrast with the suboptimality of many decision-making processes, and by the need to embrace agency with its constituents of autonomy and differentiation in all kinds of engineered systems. I especially enjoy working on problems related to societal need, problems in which the business impact is clear, and on problems that can enable significant paradigm shifts within computational systems.[2] In closing, I will mention a few examples of the kinds of problems that I would like to study in the future:

- We need methods to automatically control "anytime" optimization algorithms to make the best use of available computational resources for the instance at hand while generating a computational proof of incentive properties. While I have made some progress in this direction in the context of "anytime strategyproofness" and "computational ironing" [PS04a, PD07], much remains to be done. I believe progress will require *quantifiable measures of truthfulness*, that enable the field to move beyond the absolutism of *strategyproofness* and towards methods that are non-manipulable to varying degrees.

- MD addresses the incentive problem that occurs in environments with public (enforceable) actions and private information. There is also a *dual* problem that deals with the incentive problem that occurs in environments with private actions and varying degrees of public information. I conceptualize this problem as one of *incentive-based environment design* (IBED). Whereas the institutions of MD collect information from agents and then make and enforce actions, in IBED one instead associates rewards with (partially) observable world states and seeks to induce agents to choose to take desirable actions in this modified environment. IBED expands on *contract theory* and *principal-agent theory* of economics, and can find broad application to e-commerce and social-computing settings.

- I want to study the application of economic methodologies to promote efficient differentiation, learning, and computation *within* a single, intelligent system. Just as incentive-compatibile mechanisms provide a framework to promote socially-optimal decision making in society, I want to understand whether mechanisms can also find a role in the coordination of modular AI systems. In particular, dynamic incentive mechanisms may provide a framework for the optimal allocation of shared computational resources to competing and interdependent computational processes and offer a direction for addressing issues of meta-rationality in providing robust computational intelligence.

I expect to continue to improve connections with colleagues outside of computer science at Harvard, for example through giving guest lectures in classes, by meeting, advising, and serving on the committees of Business-Economics Ph.D. students, and by co-advising students in the Science, Technology and Management program. In building and maintaining bridges to economics, I am also fortunate in that there are many high quality workshops that span the fields, including those in the *DIMACS Special Focus on Computation and the Socio-Economic Sciences* ('04- '08), one of which I co-organized. In other community building efforts, I will now serve as one of four "area editors" to *Games and Economic Behavior* (the premiere journal in game theory), with my responsibility roughly assigned to papers in computer science, mechanism design, and related areas.

[2] I have been actively involved in discussions at the Federal Communication Commission related to the design of next generation marketplaces and I participated in a Federal Aviation Administration-funded project to explore market designs for landing slot allocation at LaGuardia airport.

References

[ACM06] Lawrence M. Ausubel, Peter Cramton, and Paul Milgrom, *The clock-proxy auction: A practical combinatorial auction design*, in Cramton et al. [CSS06].

[ACN+07] Alvin AuYoung, Brent Chun, Chaki Ng, David C. Parkes, Amin Vahdat, and Alex Snoeren, *Practical market-based resource allocation*, Tech. Report CS2007-0901, CSE, UCSD, 2007.

[AM02] Lawrence M. Ausubel and Paul Milgrom, *Ascending auctions with package bidding*, Frontiers of Theoretical Economics **1** (2002), 1–42.

[BHH+06] John Brunelle, Peter Hurst, John Huth, Laura Kang, Chaki Ng, David C. Parkes, Margo Seltzer, Jim Shank, and Saul Youssef, *Egg: An extensible and economics-inspired open grid computing platform*, GECON 2006 - Proceedings of the 3rd International Workshop on Grid Economics & Business Models (Singapore), 2006.

[BV07] Dirk Bergemann and Juuso Välimäki, *Efficient dynamic auctions*, Tech. Report Cowles Foundation Discussion Paper No. 1584, Yale University, 2007.

[CBA+05] Brent N. Chun, Philip Buonadonna, Alvin AuYoung, Chaki Ng, David C. Parkes, Jeffrey Shneidman, Alex C. Snoeren, and Amin Vahdat, *Mirage: A microeconomic resource allocation system for sensornet testbeds*, Proceedings of 2nd IEEE Workshop on Embedded Networked Sensors (EmNetsII), 2005.

[CPS06] Ruggiero Cavallo, David C. Parkes, and Satinder Singh, *Optimal coordinated learning among self-interested agents in the multi-armed bandit problem*, Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence (UAI'2006) (Cambridge, MA), 2006, pp. 55–62.

[CPS07] ______, *Efficient online mechanisms for persistent, periodically inaccessible self-interested agents*, DIMACS Workshop on The Boundary between Economic Theory and Computer Science, 2007.

[CSS06] Peter Cramton, Yoav Shoham, and Richard Steinberg (eds.), *Combinatorial auctions*, MIT Press, January 2006.

[CT05] Yan Chen and Kan Takuechi, *Multi-object auctions with package bidding: An experimental comparison of iBEA and Vickrey*, Tech. report, School of Information, University of Michigan, 2005.

[FRS06] Joan Feigenbaum, Vijay Ramachandran, and Michael Schapira, *Incentive-compatible interdomain routing*, Proc. of the 7th Conference on Electronic Commerce (EC'06), 2006, pp. 130–139.

[FS02] Joan Feigenbaum and Scott Shenker, *Distributed Algorithmic Mechanism Design: Recent Results and Future Directions*, Proceedings of the 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications, 2002, pp. 1–13.

[FSS06] Joan Feigenbaum, Michael Schapira, and Scott Shenker, *Distributed algorithmic mechanism design*, in Cramton et al. [CSS06].

[HKMP05] Mohammad T. Hajiaghayi, Robert Kleinberg, Mohammad Mahdian, and David C. Parkes, *Online auctions with re-usable goods*, Proc. ACM Conf. on Electronic Commerce, 2005, pp. 165–174.

[HKP04] Mohammad T. Hajiaghayi, Robert Kleinberg, and David C. Parkes, *Adaptive limited-supply online auctions*, Proc. ACM Conf. on Electronic Commerce, 2004, pp. 71–80.

[KP07] Laura Kang and David C. Parkes, *A decentralized auction framework to promote efficient resource allocation in open computational grids*, Proc. Joint Workshop on The Economics of Networked Systems and Incentive-Based Computing (San Diego, CA), 2007.

[LCP05] Sébastien M. Lahaie, Florin Constantin, and David C. Parkes, *More on the power of demand queries in combinatorial auctions: Learning atomic languages and handling incentives*, Proc. 19th Int. Joint Conf. on Artificial Intell. (IJCAI'05), 2005.

[LJC+07] Benjamin Lubin, Adam I. Juda, Ruggiero Cavallo, Sébastien M. Lahaie, Jeffrey Shneidman, and David C. Parkes, *ICE: An Expressive Iterative Combinatorial Exchange*, Journal of Artificial Intelligence Research (2007), Submitted, under review.

[LP04] Sébastien M Lahaie and David C Parkes, *Applying learning algorithms to preference elicitation*, Proc. ACM Conf. on Electronic Commerce, 2004, pp. 180–188.

[MP07] Debasis Mishra and David C. Parkes, *Ascending price Vickrey auctions for general valuations*, Journal of Economic Theory **132** (2007), 335–366.

[NBC+05] Chaki Ng, Philip Buonadonna, Brent N. Chun, Alex C. Snoeren, and Amin Vahdat, *Addressing strategic behavior in a deployed microeconomic resource allocator*, Proceedings of 3rd Workshop on the Economics of Peer to Peer Systems (P2PECON'05), 2005.

[NS06] Noam Nisan and Ilya Segal, *The communication requirements of efficient allocations and supporting prices*, Journal of Economic Theory **129** (2006), 192–224.

[Par01] David C. Parkes, *Iterative combinatorial auctions: Achieving economic and computational efficiency*, Ph.D. thesis, Department of Computer and Information Science, University of Pennsylvania, May 2001.

[PCE+05] David C. Parkes, Ruggiero Cavallo, Nick Elprin, Adam I. Juda, Sebastien M. Lahaie, Benjamin Lubin, Loizos Michael, Jeffrey Shneidman, and Hassan Sultan, *ICE: An iterative combinatorial exchange*, ACM Conf. on Electronic Commerce, 2005, pp. 249–258.

[PD07] David C Parkes and Quang Duong, *An ironing-based approach to adaptive online mechanism design in single-valued domains*, Proc. 22nd National Conference on Artificial Intelligence (AAAI'07), 2007, pp. 94–101.

[PFP06] Adrian Petcu, Boi Faltings, and David C. Parkes, *MDPOP: Faithful distributed implementation of efficient social choice problems*, Proc. 5th Int. Joint Conf. on Autonomous Agents and Multiagent Systems, 2006, pp. 1397–1404.

[PS03] David C. Parkes and Satinder Singh, *An MDP-based approach to Online Mechanism Design*, Proc. 17th Annual Conf. on Neural Inf. Processing Systems (NIPS'03), 2003.

[PS04a] David C. Parkes and Grant Schoenebeck, *Growrange: Anytime VCG-Based Mechanisms*, Proc. 19th National Conference on Artificial Intelligence (AAAI-04), July 2004, pp. 34–41.

[PS04b] David C. Parkes and Jeffrey Shneidman, *Distributed implementations of Vickrey-Clarke-Groves mechanisms*, Proc. 3rd Int. Joint Conf. on Autonomous Agents and Multi Agent Systems, 2004, pp. 261–268.

[PSY04] David C. Parkes, Satinder Singh, and Dimah Yanovsky, *Approximately efficient online mechanism design*, Proc. 18th Annual Conf. on Neural Inf. Processing Systems (NIPS'04), 2004.

[SNP+05] Jeffrey Shneidman, Chaki Ng, David C. Parkes, Alvin AuYoung, Alex C. Snoeren, Amin Vahdat, and Brent Chun, *Why markets could (but don't currently) solve resource allocation problems in systems*, Proceedings of Tenth Workshop on Hot Topics in Operating Systems) [HotOS-X 2005, 2005.

[SP04] Jeffrey Shneidman and David C. Parkes, *Specification faithfulness in networks with rational nodes*, Proc. 23rd ACM Symp. on Principles of Distributed Computing (PODC'04) (St. John's, Canada), 2004, pp. 88–97.

[SPM04] Jeffrey Shneidman, David C. Parkes, and Laurent Massoulie, *Faithfulness in Internet algorithms*, Proc. SIGCOMM Workshop on Practice and Theory of Incentives and Game Theory in Networked Systems (PINS'04) (Portland, USA), 2004.

[SSB07] Stefan Schneider, Pavlo Shabalin, and Martin Bichler, *On the robustness of primal-dual iterative combinatorial auctions*, Tech. report, Dept. of Informatics, TU München, Germany, 2007.