Behavior, Computation and Networks in Human Subject Experimentation

NetSE Workshop Report
Michael Kearns (Computer and Information Science, University of Pennsylvania)
Colin Camerer (Economics, Caltech)

Background: *On July 31 and August 1, 2008 we held a NetSE workshop on the topic of "Behavior, Computation and Networks in Human Subject Experimentation". This interdisciplinary workshop was designed to bring together a relatively small number of researchers with the following broad profiles:*

- *Researchers from economics, game theory and sociology whose interests include behavioral human-subject experiments.*
- *Researchers from economics, game theory and sociology interested in computational and algorithmic models for behavior, and algorithmic issues more broadly (such as equilibrium computation).*
- *Researchers from computer science with interests in game theory, economics and sociology, especially behavioral, experimental and simulation work.*
- *Researchers from all of the fields above with interests in social, organizational, technological and other networks, and how network structure and formation interact with individual and collective behavior.*

The remainder of this document describes the scientific agenda and research challenge areas emerging from discussions at the workshop and beyond.

1. Rationale for a New Research Agenda

Researchers in the computer science, economics, game theory and sociology communities have been engaged for some time now in healthy and vibrant interaction on a variety of *theoretical* topics. We assert that the natural and most important next frontier in this dialogue is the introduction of a *behavioral and experimental* component. Of particular interest are organizations and systems in which an underlying network structure strongly governs interaction and strategy.

Regarding the dialogue so far, between economics and computer science we have the well-established field of *algorithmic game theory*, whose work can now be found in multiple journals and conferences of both fields (ACM Conference on Electronic Commerce; Workshop on Internet Economics (WINE); STOC, FOCS and SODA of the theoretical computer science community; World Congress of Game Theory; Games and Economic Behavior; and many other examples); in an extensive recent edited volume from Cambridge University Press; in many notable publications co-authored by members of both communities; and so on. Similarly, by now there is a fair amount of contact between primarily mathematical topics within sociology (such as the diffusion of trends within a social network) and the theoretical computer science community.

If we are to take such interactions as more than theory for its own sake --- by which we mean that they might provide the foundation for an empirical science that is applicable to real problems and data, and able to make predictions (and potentially policy recommendations) --- then it is clear that we must begin to develop a heretofore missing behavioral and experimental component. In the same way that behavioral game theory and economics seek to adapt their theoretical counterparts towards actual human and organizational behavior (thus improving their applicability), we seek to build a behavioral and experimental discipline encompassing strategic settings important in computer science and technology, network science, and related fields.

As we shall discuss in the following section, building this discipline presents a number of significant conceptual, scientific and resource challenges to the constituent research communities. We believe the reward for meeting these challenges will be the creation of an important new field whose content will be widely applicable to the myriad modern problems in which strategic considerations, technology and behavior interact.

2. Emerging Research Challenges

Unifying Algorithmic and Behavioral Game Theory

In their own fashions, both algorithmic and behavioral game theory seek to "repair" classical game theory, arguably in the direction of "realism". Ideally game theory and related fields would provide accurate predictions of actual strategic behavior in individuals and organizations. Behavioral game theory seeks to reconcile theoretical models with empirically observed behavior in controlled experiments. Algorithmic game theory attempts to identify and rectify classical equilibrium notions by enforcing plausible demands on computational and other resources.

Ideally these two approaches should be unified and refined --- behavioral models taking more precise account of computational considerations, and algorithmic models evaluated and improved in light of experimental evidence. The mathematical, methodological and cultural chasms between the two communities are large, which is a significant part of the challenge. Algorithmic models will need to be refined in ways less related to traditional computational complexity (P vs. NP, and the various subclasses of P) and more related to cognition. But there is already promising work in this direction that crosses the disciplinary boundaries --- for example, recent work of Camerer and colleagues on "cognitive hierarchies" of varying levels of strategic behavior that directly account for computational limitations. We believe both algorithmic and behavioral approaches to game theory are sufficiently mature independently that the attempt to build a unified theory has arrived.

Network and Systems Infrastructure for Behavioral Experiments

At both the workshop and during a long series of dialogues between several participants, there has been discussion and excitement around the possibilities of building internationally shared networking and systems infrastructure for the conducting and

support of large-scale behavioral experiments in sociology, game theory, economics, and most recently, computer science. The fundamental observation is that the Internet, Web and other technologies have created the possibilities for (semi-)controlled experiments in these disciplines --- where small population sizes and the difficulties of human subject management have long been limiting factors --- on a large or massive scale. There are already numerous examples of such experiments, but each has employed highly specialized software and technology.

A significant portion of workshop time was devoted to discussing what such shared infrastructure might provide, how general it should be, what precedents there are, whether existing commercial platforms (such as Amazon Mechanical Turk) might suffice or at least serve as models, and many other issues. Perhaps the most important design issue is the trade-off between generality and ease of use, while the most important methodological concern is the maximal retention of experimental control and subject knowledge and management.

There is a strong sense among the stakeholders that (a) designed and implemented properly, such an experimental platform could have a transformative effect on behavioral research, and (b) its creation would be an extremely challenging and resource-intensive project requiring close collaboration between sociologists, economists, and computer scientists.

Theory and Design Principles for Peer Production

Related but distinct from the topic of shared experimental infrastructure is the phenomenon of recent "human peer production" --- systems in which massive numbers of distributed individuals voluntarily "solve" collective problems or build influential and useful artifacts. The diversity of such systems across several dimensions (problem solved, nature of individual contributions, incentives, etc.) is staggering and includes Wikipedia, the ESP Game, del.icio.us, Amazon Mechanical Turk, NASA click workers, Galaxy Zoo, prediction markets of many kinds, social networks, and many other examples. Yet there is essentially no theory about the design or behavior of such systems, including on basic questions such as contributor population size vs. quality of collective outputs, choice of incentive schemes, dealing with problems not easily decomposed into "modular" subtasks, and so on. The moment seems right to tackle such challenges with an interdisciplinary approach. (Duncan Watts and Kearns have held a series of DARPA workshops on this broad topic, with participation from many at the NetSE workshop as well as from industry and the military.)

3. Workshop Format

The workshop participants numbered approximately 20 active scientists representing multiple disciplines, each of them highly influential and at the forefront of their respective fields. We felt that assembling a group with this "gravitas" was important both

to obtain mature and accurate representations of the interests of the constituent areas, but also for follow-up evangelization of our nascent and emerging research agenda.

The workshop began with brief research presentations from each scientist. While the goal was to maximize the time allotted to semi-structured discussion and brainstorming, the diversity of backgrounds, interests and terminology was sufficient to warrant steeping the participants in work and viewpoints of the others. These talks were by themselves fascinating for both their individual content and collective variety and interconnections.

The bulk of the time, however, was devoted to informal and open discussion on a small number of central topics. The complete agenda, with participants, talk titles and discussion topics is provided in the following section.

4. Workshop Agenda and Participants

Behavior, Computation and Networks in Human Subject Experimentation
Thursday, July 31 and Friday, August 1, Del Mar Inn, Del Mar CA

Agenda

Thursday, July 31

9:00 Welcoming remarks: M. Kearns (Penn) and C. Camerer (Caltech)
9:15 Remarks from E. Zegura (Georgia Tech) on NetSE Council interest in our topics
9:30 Brief introductions
9:45 Brief research/project presentations (~15 minutes each):

V. Crawford (UCSD): *Studying Strategic Thinking by Monitoring Search for Hidden Payoff Information and Analyzing the Data in the Light of Algorithms that Link Cognition, Search, and Decisions*

M. Wellman (Michigan): *Software Agents and Empirical Game Analysis*

J. Ledyard (Caltech): *Agent-based models for repeated game experiments*

Break

C. Camerer (Caltech): *Evidence of algorithmic game theory from human experiments*

M. Kearns/S. Judd (Penn): *Behavioral Network Science and the Democratic Primary Game*

M. McCubbins/M. Paturi/N. Weller (UCSD): *Effects of Complexity, Incentives and Network Structure on Multi-Player Coordination Games*

J. Fowler (UCSD): *Eat, Drink, and Be Merry: The Spread of Obesity, Substance Use, and Happiness in a Large Social Network*

12:30 Lunch

1:30 Brief research presentations, continued:

S. Kariv (Berkeley): *A Normal Form Game Experiment of Trading Networks*

A. Pfeffer/K. Gal (Harvard): *Modeling the reasoning of people and computer agents in strategic settings*

B. Rogers (Northwestern): *Communication Networks: An Experimental Study of Influence*

2:30 Discussion of the morning's presentations: common themes and differences; marrying different approaches; what's missing; etc.
3:00 Brainstorming Topic 1: Algorithmic game theory and behavioral game theory/economics
4:00 Break
4:15 Brainstorming Topic 2: Relevance/incorporation of simulated agents in behavioral experiments
5:00 Adjourn
6:45 Informal dinner overlooking the Pacific, Martin Johnson House, Scripps Institute of Oceanography, La Jolla

Friday, August 1

9:00 Brief research presentation by D. Watts, Yahoo! Research/Columbia: *Virtual Labs: Using the Web to Conduct Human Subjects Experiments*

9:15 Recap of Thursday, discussion of new topics
9:30 Brainstorming Topic 3: "Scaling Up" behavioral experiments: use of the web, Amazon Mechanical Turk, peer production, etc. Do we need a "programmable infrastructure"?
10:30 Break
11:00 Brainstorming Topic 4: What are the applications of all this stuff?
12:30 Lunch
2:00 Brainstorming Topic 5: Where do we go from here?