D. Project Description

D.1 Introduction

D.1.1 Motivation

The global economy is in the midst of large-scale transformation. Innovations in information technology (IT) – the explosive growth of the Internet, of capacities for storage and data transmission, of wireless connectivity, etc – are reshaping the relations among business firms and consumers. In turn, technological innovation is shaped by business realities. The transformation is still in its early stages, but has already produced unexpected cycles of resistance and adaptation, as witnessed in the dot.com bubble and crash, the recent waves of outsourcing and off-shoring, and the rise of the new service economy.

Services now constitute over 2/3 of the US economy; up from less than half only a few decades ago (Smith, Gruenberg, Harris, and Strassner, 2005). In particular, US job growth in knowledge-intensive business services continues to expand while manufacturing job loss widens (Porat, 1977; IFTFR, 2004). Likewise, the US is enjoying a growing trade surplus in services while the merchandise trade deficit widens. Computer hardware industry leaders like IBM and HP have transformed themselves into services companies; indeed, over half of IBM's revenue is now from services. By now most consumers in first-world countries have Internet access supporting electronic B2C (business to consumer), C2C markets, and B2B exchange (e.g., http://www.worldwideretailexchange.org); dollar volume now is in the trillions.

We are now at a crucial juncture. The naïve enthusiasm of the dot.com bubble has evaporated, and it is now clear to everyone that the laws of economics have not been repealed. But precisely how do the laws play out in the transformed economy? Research lags as the transformation re-accelerates. Many specific questions demand answers. For example:

- B2C and B2B Internet markets currently use traditional market formats such as Posted Offer and Call Market, described below (e.g., by Yahoo Auctions, and Fastparts.com), along with new variants and combinations (e.g., eBay's Buy It Now option). How profitable and socially efficient are the current (and as-yet untried) market formats in the non-traditional IT environments?
- The costs for buyers and sellers to enter or exit a market or switch transaction partners are lower in Internet markets than in bricks-and-mortar markets, creating an incentive for more volatile market participation. What are the economic and social implications?
- Automated agents are increasingly used to buy and to sell on the Internet. What are the implications of agent-human interactions?[1]
- A second wave of e-commerce features the creation and sale of new composite products such as wireless connectivity, "utility computing," "computons" and "business on demand." What sorts of market formats can best cope with such customized services?
- Buyers and sellers may have concerns other than immediate self-interest, e.g., a concern for reciprocity and trust. How should these considerations be taken into account when designing automated agents and market formats?

We propose an integrated series of research projects designed to answer such questions. The multidisciplinary research team brings together economists and computer scientists experienced in relevant research and in close contact with practitioners in the emerging economy. The main research tools are theoretical models (formal analytic models as well as simulations), laboratory experiments with human subjects, and Internet field trials. The work emphasizes novel features of IT environments, viz:

- asynchronous participation, or, in economic jargon, endogenous entry and exit by sellers and buyers;
- automated agents, which interact in real-time markets with humans and with other agents;
- new market formats enabled by IT; and
- new composite services with features that are individually customized in each transaction.

[1] An early but cautionary example is the US stock market collapse of October 1987. Analysts (e.g., Brady et al, 1988) emphasize the role of automated agents executing sell orders for "portfolio insurance." Price continuity, a key assumption of portfolio insurance, was broken at the crucial moment by the automated agents that relied on it.

D.1.2 Objectives

The overall objective is to identify the market formats and agent behaviors that are most effective in realistic IT environments. Measures of effectiveness include the traditional economic criterion of static (or equilibrium) efficiency: individual seller profits, individual consumer surplus, and aggregate gains from trade. We will also use two less traditional measures of effectiveness. Dynamic (or learning) efficiency considers the gains or losses accrued while the system converges to equilibrium, while evolutionary stability considers the ability of one market format to attract buyers and sellers from rival formats. Our concern always is to spread the social benefits of the economic transformation and help prevent unnecessary economic and social disruption.

The proposed work has three major components. The first examines one-sided market (or auction) formats. The Economic Science Laboratory (ESL) at the University of Arizona will take the lead here. The second component is two-sided markets, led at the Learning and Experimental Economics Projects (LEEPS) Lab at the University of California, Santa Cruz. The third component is composite service markets led at the Almaden Service Research (ASR) group at the IBM Almaden Research Center in San Jose, California. Specific objectives for each component project are noted below.

D.1.3. Expected Significance and Broader Impact

The proposed research falls mainly in the Agents of Change (AOC) area of emphasis but also has implications for the Dynamics of Human Behavior (DHB) and the Decision Making, Risk and Uncertainty (DRU) areas. Currently, the US economy is still recovering from a recession triggered by the dot.com crash. Many observers believe that the crash was the result of commercial hyperbole running too far ahead of science.[2] Properly incorporating the new IT into the economy, transforming markets and services, is a key to economic growth and greater social equity. The transformation itself is largely in the hands of business people and policy makers, and their work will be greatly aided by science. The IT science is solid but the economic science lags. We propose fundamental research to help close the gap.

We shall develop new concepts, theories and models for IT-enabled markets, and will test and refine them in computer simulations and in laboratory experiments with human subjects and automated agents. We will also begin to explore Internet field experiments and tournaments. Markets once were mainly local, but now increasingly coordinate and control human activities on a national and global scale, echoing themes of the DHB area of emphasis. And innovations in market strategies certainly reflect the main DRU themes at both the individual and group level.

The proposed research is fundamental science, but the PIs' ties to e-commerce professionals will help integrate the findings quickly into commercial practice. Two of the labs are located in the orbit of Silicon Valley, the international center of high technology and e-commerce, and the third is located in Tucson, a leading regional center; see section D.4.3 below for supporting material on lab resources. Practitioners at large, established enterprises (e.g., IBM and HP) as well as numerous local startups will have access to the intelligent trading agents we propose to develop. We anticipate that the market formats we develop, especially for composite services, will prove even more useful in practice. The platforms we develop for conducting and testing alternative formats may also turn out to have practical value.

The broader impact definitely includes education. We shall make the results and the software available to all research and teaching colleges via EconPort (http://www.econport.org), an economics component of the National Science Digital Library (http://www.nsdl.org) that is developed and maintained by ESL. Students in IT-oriented programs as well as students in economics departments and business schools will be able to learn how to participate in electronic markets with other humans and with trading agents.

[2] The argument is as follows: Spurred by pundits who claimed that the new IT technology somehow suspended the laws of economics, investors poured money into venture capital firms, which channeled the funds into all sorts of IT projects. Many of the projects had dubious prospects for attracting customers (much less for earning profits), inflating a bubble that could only (in hindsight) end in a crash.

Closer to home, graduate students from our two host universities will advance their educational goals by serving as research assistants. It is worth mentioning that underrepresented minority groups have had especially good opportunities at these universities and our laboratories. Located in the Southwest, the University of Arizona has an unusually high percentage of Hispanic and Native American students. Both current postdoctoral research associates at ESL are women, and many of the graduate research assistants are members of minority groups. The Santa Cruz campus of the University of California has an even higher percentage of students classified as Hispanic, Black, and Asian, and the Economics program attracts a disproportionate share of them. In recent years, more than half of the graduate research assistants at LEEPS have been women or members of minority groups.

D.2 Present State of Knowledge

Our proposed work builds on research from several disciplines, which we shall now summarize briefly. HSD reviewers are quite heterogeneous and each is likely to find that too much space is devoted to some lines of research and not enough to others. If reviewers disagree on which lines should be covered more fully or more briefly, perhaps we have done our job properly!

D.2.1 Experimental Economics Techniques and Auction Formats

During the last two decades, experimental methods became integrated into mainstream economics, a methodological advance celebrated by the 2002 Nobel Prize in Economic Sciences. The methods for conducting laboratory experiments can be explicated with an example from single-item auctions.

The testable theory originated in a classic paper by Vickrey (1961). He analyzed four auction formats: (1) the ascending-price, real-time (English) auction in which the bidder with the last (i.e., highest) bid buys the auctioned item at a price equal to his/her bid; (2) the descending-price, real-time (Dutch) auction in which the bidder with the first bid buys the auctioned item at a price equal to the bid; (3) the first-price sealed-bid (non-real-time) auction in which the bidder with the highest bid buys the auctioned item at a price equal to his bid; and (4) the second-price sealed-bid (non-real-time) auction in which the bidder with the highest bid buys the auctioned item at a price equal to the second-highest bid.

Vickrey developed the Bayesian-Nash equilibrium analysis for these four auctions: each bidder adopts a bid function that is the best strategic reply to the bid functions adopted by rival bidders. The theory yields predictions about the bid functions, and the resulting probability distribution of seller's revenue, in each of the four auction formats. An aggregate prediction is "revenue equivalence": if bidders are risk neutral, then all four formats generate the same expected revenue. Furthermore, if bidders are risk averse then the first-price and Dutch auctions are predicted to yield the same expected revenue, higher than the common expected revenue from the English and second-price sealed-bid auctions.

The predictions can be tested in a laboratory experiment, as follows. Vickrey's theory assumes that the bidders' values for the auctioned item are independently drawn, from a distribution known to everyone, and that each bidder knows the actual value he drew but not his rivals' values. This informational assumption can be implemented directly in the laboratory. The experimenter simply chooses and announces the distribution, e.g., uniform on [0, 100], and privately draws values for each bidder, say v_i, for $i = 1, 2, ..., n$. Economic incentives are induced on the subject bidders (Smith, 1976) by the experimenter's promise to pay the high bidder an amount of money equal to v_i less the amount of her bid (or the second-highest bid, in the case of the second-price auction). After running a series of auctions for each treatment, the subjects' bids and induced values for the abstract auctioned item can be used to test the predictions of the theory. For example, individual-subject bid and value data can be used to econometrically estimate the parameters of individual bid functions and test the parameter restrictions of the theory. Individual subject data can be used to test whether bidding behavior changes in the predicted ways as the number of bidders is changed and as the type of auction is changed. Market price data can be used to test revenue equivalence.

Cox, Roberson, and Smith (1982) conducted laboratory tests of the theory and report: (a) bids and prices systematically higher than predicted in first-price sealed-bid auctions; (b) bids systematically lower

in the Dutch auction than in the first-price auction; and consequently (c) violations of revenue equivalence. These findings led to the development and testing of new theoretical models (e.g., Cox, Roberson, and Smith, 1982; Cox, Smith, and Walker, 1983, 1988; Cox and Oaxaca, 1996).

Laboratory evidence is complementary with field evidence, whether the field evidence is collected passively from naturally-occurring markets or from (partially) controlled field experiments. A disadvantage of field evidence is that it does not permit observation (much less control) of bidders' values for the auctioned item. Hence the predictions regarding bid functions cannot be tested. Naturally-occurring field data permit tests of revenue equivalence, but the tests often are inconclusive because the choice of auction format may co-vary with other bid-determining features of the environment. Field experiments, however, can control the auction format. Lucking-Reiley (1999) employed such controls in an Internet (field) experiment and also found that revenue equivalence for Dutch and first-price auctions fails empirically but that Dutch auction prices were *higher*. The apparent inconsistency with the earlier Cox, Roberson and Smith finding of *lower* prices in the Dutch auction was explained by Katok and Kwasnica (2002) as the result of different Dutch auction (price) clock speeds. Cox, Roberson, and Smith used a fast clock speed similar to bricks-and-mortar Dutch auctions involving professional bidders, such as the Aalsmeer Flower Auction (http://www.vba.nl), whereas Lucking-Reiley used a very slow clock speed considered more likely to typify Internet auctions with amateur (consumer) bidders. This example illustrates the general point that researchers need to be sensitive to changes in auction formats that are introduced when they are moved to the Internet: seemingly small changes in market procedures sometimes can fundamentally change market performance.

D.2.2 Agent Technology

Agent technology is one of the fastest growing areas of research and system development in Information Technology (Jennings, Sycara, and Wooldridge, 1998; Weiss, 2000). It represents a new way of modeling complex information management and decision tasks. Agents also represent a new computing environment for designing and implementing complex software systems (Joshi and Singh, 1999). Agents have begun to attract the attention of economists (e.g., Tesfatsion, 2004) under the rubric of agent-based computational economics.

In this proposal, the term *agent* is shorthand for automated agent or trading agent, and refers to a computational entity, a computer algorithm. More specifically, an agent is an automated strategic decision-making and execution system that operates in a networked environment on tasks delegated by their human users. Agents interact with other agents or with human players directly or indirectly through specified online market formats. Much of the existing agent work in IT focuses on information management tasks associated with the Internet (Chau, Zeng, Chen, and Huang, 2003; Sycara and Zeng, 1996; Jennings, Sycara, and Wooldridge, 1998). Well-known examples include systems that provide strategic decision support in international trade negotiation, labor negotiation, bargaining, and online auctions (Kraus and Wilkenfeld, 1993; Zlotkin, 1994; Zeng and Sycara, 1998; Parkes, Ungar, and Foster, 1999; Sandholm, 2000; and Greenwald and Stone, 2001). Most of these current systems are programmed with rules-of-thumb distilled from intuition and good behavioral practice in human strategic interactions.

Computational models of auctions and bargaining from the Economics and Artificial Intelligence (AI) literature will serve as the basis for constructing such agents. These include optimization-based strategies for competitive bidding that assume known distributions of values and the number of bidders (L. Friedman, 1956), AI logic and rule-based negotiation strategies (Kraus, Sycara, and Evanchik, 1998; Kraus, Wilkenfeld and Zlotkin, 1995), automata play (Rubinstein, 1986, 1998), and adaptive auction and negotiation strategies using machine learning techniques (Mor, Goldman and Rosenschein, 1995; Wellman and Hu, 1998; Zeng and Sycara, 1998).

D.2.3 Two-sided Market Experiments, Agents and Tournaments

The literature on laboratory market experiments with human traders goes back to Smith (1962); see Smith (1982), Plott (1982) and Davis and Holt (1993, chapters 3-4) for surveys. The experimenter can control directly the market format as well as the structural features, e.g., the numbers of buyers and sellers and

their values and costs. The stylized fact is that, given an effective market format such as the continuous double auction (CDA) described in section D.3.2 below, laboratory markets quickly converge to the efficient competitive equilibrium for a wide range of value and cost structures, even with as few as three buyers and three sellers. Of course, this result is not yet confirmed in Internet-like environments with switching costs and asynchronous entry and exit.

Occasional studies going back to (J.) Friedman and Hoggatt (1980) used simple trading agents in the same market as humans. The original motivation was to simulate demand in oligopoly settings, but studies such as Cason and (D.) Friedman (1997) use agents to provide a more stable structure for learning equilibrium behavior. Several recent laboratory asset market studies (e.g., Plat, 1995) use simple money-losing agents as "noise traders" to encourage human traders to become active. In an influential article, Gode and Sunder (1993) use simple "zero-intelligence" (ZI) agents to demonstrate rapid convergence to competitive equilibrium in the CDA even without rational traders. However (as noted by Cason and Friedman), the ZI agents create much less efficient outcomes than humans in other market formats such as the call market (CM), also described in section D.3.2. As for the interactions of ZI agents with humans, unpublished work by Sunder and others suggests that ZI agents can exploit human traders (by virtue of their greater speed) in the CDA but are easily exploited by humans in the CM. Selten and Buchta (1998) use human controlled agents (referred to as the strategy method) to study convergence to Bayesian-Nash equilibrium bidding strategies in single-sided auctions.

The standard reference on CDA and other two-sided market formats, Friedman and Rust (1993), surveys theoretical analysis and empirics available at that time. Recent work studying agents more sophisticated than ZI in two-sided markets includes Gjerstad and Dickhaut (1998), Das, et al. (2001), Zhan et al (2002), Cliff (2003), and Gjerstad (2005). Their agents are in principle exploitable by experienced human traders, but several of the studies show that in practice the agents can outperform inexperienced humans.

An instructive study of human controlled agents is reported in Rust, Miller and Palmer (1993), who analyze the results of a market tournament sponsored by the Santa Fe Institute. Smart agents designed by contestants competed in a newly created market format called a synchronized double auction (SDA). The winner was a strategic but relatively simple agent that waits in the background while others negotiate and then jumps in at the last minute to "steal the deal," i.e., what now is called a sniping agent. However, as such agents gain market share, their profitability plummets and so does market efficiency. Rust, et al. interpret this finding as an unstable market ecology or, in our terminology, ineffective learning in the SDA. A careful reading of their study suggests an alternative interpretation. The strict separation of "buy/sell" steps from "bid/ask" steps in the SDA allows the winning agent to exploit the information revelation of more forthright agents. Integration of these steps, as in the standard continuous double auction (CDA) market format, apparently would disable the parasitic agent and restore efficiency. It also seems clear that the SDA is evolutionarily unstable: exploited traders would move from the SDA to a CDA where their deals couldn't easily be stolen. The moral of this story is that effectiveness (static efficiency as well as learning efficiency and evolutionary stability) can turn on seemingly minor details of the market format. Careful theoretical and empirical studies are essential to understanding and predicting market outcomes.

D.2.4 Service Markets and Agent Technology

Research on the growing services economy has begun in earnest[3] given the large-scale economic changes that are becoming increasingly evident. The complexity of services has been a fundamental reason why

[3] Workshop on Service Innovations for the 21st Century
(http://www.almaden.ibm.com/asr/events/serviceinnovation/)

there has been no laboratory research on the subject in economics. Services typically include multiple dimensions and dynamic coordination which contrasts sharply with the simple and identical units of exchange found in idealizations of economic theory and laboratory experiments. The basis for much of fundamental economic theory focuses on the distribution and exchange of commodities and manufactured products (Smith, 1904; Shaw, 1912). Historically, services were thought of as just some sort of intangible residual in the process of selling and delivering goods, adding overhead and expense (Rathmell, 1966).

As more and more companies face the commoditization of their product markets, they have been force to add value to offerings by changing to a service model (e.g. provide product in the context of a service). As such, a greater understanding of services in practice has been emerging. Most would agree now on the follow definition: Services are activities that produce value through the coordination of skills, knowledge, and experience in the context of an ongoing relationship between the parties involved. This implies that both the client and the provider have responsibilities that they must fulfill to realize the value of the service (Vargo and Lusch, 2004). The buyer must work with the seller (i.e., coordinate) by monitoring their business, operations, and the service being provided and communicating needs, requirements, and goals. Sometimes this means negotiating ongoing service arrangements with the seller in order to better fit the service to the buyer's business or, more generally, needs. The better the coordination, the greater the value produced. Coordination, however, is not always required to produce value for certain services (e.g., doctors). Services may be composite and customized offerings made up of knowledge, skills, capabilities, innovations, composite and customized goods, and commodities (Rust, 1998). Many companies, particularly in the IT industry, have referred to services as *solutions*. While services seek to maximize the value of the activity, solutions provide a way to solve a specific problem. Solutions, however, provide a more comprehensive fix and better guarantees than selling a product alone.

The lack of markets for services is a sign of the complexity involved in buying and selling services. While agent technology provides opportunities for better management of complex interactions, negotiations, and fast-paced electronic markets, agents can, like other computer interfaces, impose costs. Even though these costs have not been systematically studied in the domain of human-supervised agent markets, they can significantly reduce both the utility of market participation as well as decrease dynamic market efficiency. Human-agent interfaces fall into three general categories of interaction style: (a) Configured agents, (b) Attentive agents, and (c) Adaptive learning agents.

Configured agents are agents whose functionality has already been implemented while more detailed behavior is configured with the expertise and interests of an individual or organization. Configuration is achieved by transferring expertise and interests to the agent using some language or representation. The language can be a set of rules, policies (Campbell, Kandogan, November, Barrett, and Maglio, 2005), rankings, or feature weightings. The representation can be text, any one of various pseudo-scripting languages, or even as graphical flow-charts (Blackwell, von Konsky, and Robey, 2001). The dynamics of configuration can differ as well with some systems requiring a complete specification before deployment (Uszok, Bradshaw, Johnson, Jeffers, Tate, Dalton, Aitken, 2004), and other systems allowing for interactive specification via elicitation (Conen and Sandholm, 2001, 2002) or Wizards (Tuchinda and Knoblock, 2004). Configured agents give buyers and sellers greater control over behavior, but typically this requires more time and skill to configure the agent. Attentive agents, alternatively, attend to user actions in combination with various learning algorithms to infer user expertise and interests (Maglio & Campbell, 2003). One of the goals is to reduce the costs of manually specifying this information. A drawback of this approach is that it generally takes a long period of time and many examples before the agent has a reasonably complete representation of user interests. Even so, the agent may never sample boundary conditions and thus, may lack controls during the most critical events. Adaptive learning agents can employ a range of statistical and machine learning methods to model and optimize a set of environmental variables (Faratin, Lee, Wroclawski, 2003). These types of agents have the highest potential to reduce use-costs and can response more dynamically to changing events but tend to require more standardized markets---markets that have a constrained set of product features, clear pricings per feature and per configuration, and preferences that can be modeled as weightings on the feature space.

As the intelligence and adaptability of agents increase, so to does their complexity, leading to potential problems with situation awareness and error recovery (Endsley, 1996; Bailey, Etgen, and Freeman, 2003). The problem has also been termed the "out-of-the-loop" problem (Endsley, 1995) and is associated with research in the automated control of commercial aircraft, air traffic controls, and nuclear power plant operators to name a few. The problem can be particularly bad for agents that have complex rules and dynamics but provide little feedback. Agents such as adaptive learning agents can be even more difficult to understand and predict. This understanding of agent capabilities and performance through feedback has been termed agent transparency (Dowling, 2000). As agent transparency decreases, users' acceptance and trust decreases (Chavez, Dreilinger, Guttman, and Maes, 1997). This is true particularly for agents that are performing substantial monetary transactions (Guttman, Moukas, and Maes, 1998). The creators of Kasbah (Chavez and Maes, 1996), a web-based C2C system in which customers created buying and selling agents for bilateral negotiation (BLN) or haggle, found a creative way to overcome the problem of agent transparency. They constrained the buyers with one of three negotiation strategies: anxious, cool-headed, and frugal, that mapped to bidding functions: linear, quadratic, or exponential. While these may not be the most optimal strategies, they were intuitive enough to help users understand how the agents would act in the marketplace and deploy them accordingly.

Careful consideration of how people, agents, and technology are arrange and coordinated is essential for designing more efficient markets and opening the door for the service markets of the future. Studies mentioned in the following sections will be conducted in such a manner as to measure performance costs, effort, and expectations of users attempting to use new agent technology in IT-enabled and service markets.

D.2.5 Behavioral Considerations: Bounded rationality, Learning and Social Preferences.

The 2002 Nobel Prize in Economic Sciences was shared by psychologist Daniel Kahneman, who used laboratory experiments to document that people depart systematically from rational self-interest in a number of environments. Departures from rationality occur when people use sub-optimal heuristics to make choices, or have not yet fully adapted to or learned about the choice environment (e.g., Cheung and Friedman, 1997; Cox, Shachat, and Walker, 2001). Camerer, et al. (1993) introduce a new technology, called mouselab, for inferring human subjects' reasoning process by observing which boxes they click to acquire information prior to making decisions in games. See also Costa-Gomez, et al. (1998).

A different sort of departure occurs when people have concerns beyond self-interest, such as loyalty to cultural norms, altruistic concerns, or a taste for positive (or negative) reciprocity. For example, a bargainer may become so outraged at the other person's "unfair" behavior that he refuses a mutually advantageous offer (e.g., Güth, et al., 1982; Cox, Friedman, and Gjerstad, 2004). Both sorts of departure may arise in our research as we look at richer laboratory and field environments, especially those involving bargaining and agents tuned by humans. We shall explain how such behavioral considerations enter into our research with software agents as we outline the plan of work below.

D.2.6. Results from Prior NSF Support

Items A and C below refer to current support rather than prior support. They are included due to close connections to the proposed work.

A. Cox, PI; Chen, Zeng, co-PIs, "An Active Object-Based Digital Library for Microeconomics Education," Oct 2002-April 2005. (NSF DUE-0226344, $700,000). This project developed EconPort (http://www.econport.org), an economics component of the National Science Digital Library (http://www.nsdl.org). EconPort will contribute to the proposed project by providing a digital library implementation infrastructure for dissemination of research output into teaching and e-commerce practice. Econport will also provide some of the software needed for running experiments.

Presentations included the following workshops and professional society meetings:

- December 2002 (several team members): NSF/NSDL All Projects Meeting, Washington

- May 2003 (D. Zeng): ACM/IEEE Joint Conferences on Digital Libraries, Houston; (J. Cox): NSF Workshop on Classroom Experiments, Tucson; (J. Cox): keynote speech about EconPort at the Seventh Experimental Economics Conference of Japan, Kyoto
- June 2003 (several team members): an entire session devoted to EconPort at the International Meeting of the Economic Science Association, Pittsburgh
- October 2003 (D. Zeng): NSF/NSDL All Projects Meeting, Washington
- November 2003 (J. Cox): workshop presentation about EconPort at Renmin University, Beijing; (J. Cox): keynote speech about EconPort at the Fourth National Conference of Development of Economics and Management Laboratories, Nanjing
- December 2003 (J. Cox): NSF Joint Principal Investigators Meeting, Washington;
- (J. Cox): Second Asian Conference on Experimental Business Research, Hong Kong
- January 2004 (J. Cox): annual meeting of the American Economic Association, San Diego
- January 2005 (J. Cox, T. Swarthout, and R. Kalla): entire session at the annual meeting of the American Economic Association, Philadelphia

Publications to date include:

- D. Zeng, J. Cox, and M. Dror , "Coordination of Purchasing and Bidding Activities Across Markets," *Proceedings of the Hawai'i International Conference on System Sciences* (HICSS-37), Big Island, Hawaii, January 5-8, 2004
- D. Zeng, J. Cox, and M. Dror (2004), "Coordination of Purchasing and Bidding Activities Across Posted Offer and Auction Markets," *Journal of Information Systems and e-Business Management,* forthcoming.

B. Cox, PI: "Collaborative Research: Group Decision-Making in Strategic Market Games," 1999 - 2002 (NSF SES-9818561, $167,000). This project is related to the proposed project because it includes research on theory and experiments with endogenous entry in auctions, trusting and reciprocal behavior, and combinatorial allocations in networks. Results were presented at numerous professional conferences. Published papers include:

- J. Cox, et al., "Endogenous Entry and Exit in Common Value Auctions," *Experimental Economics*, 4(2), 2001, pp. 163-181.
- J. Cox, "Trust, Reciprocity, and Other-Regarding Preferences: Groups vs. Individuals and Males vs. Females," in Rami Zwick and Amnon Rapoport, (eds.), *Advances in Experimental Business Research*, Kluwer Academic Publishers, 2002.
- J. Cox, et al., "Competition For vs. On the Rails: A Laboratory Experiment," *International Economic Review*, 43(3), August 2002, pp. 709-736.
- D. Chakravarti, et al., "Auctions: Research Opportunities in Marketing," *Marketing Letters*, 13(3), August 2002, pp. 281-296.
- J. Cox, "How to Identify Trust and Reciprocity," *Games and Economic Behavior,* 46(2), 2004, pp. 260-281.
- J. Cox & S. Hayne, "When Does Free Riding Promote Rational Bidding?", in Rami Zwick and Amnon Rapoport, (eds.), *Experimental Business Research*, vol. II, Kluwer Academic Publishers, forthcoming.
- J. Cox & C. Deck, "On the Nature of Reciprocal Motives," *Economic Inquiry*, forthcoming.

C. Friedman, PI; Huberman, co-PI: “Market Dynamics in Cyberspace,” 2000-2005 (NSF IIS-9986651, $435,000). This ongoing project investigates human and agent behavior in several asynchronous laboratory environments, and in Internet field data and simulations. It does not include market format comparisons, but does include development of Java-based software for a prototype CDA market. The work has been presented at more than a dozen professional workshops and conferences and has led to nine papers so far. Those most relevant to the proposed research include:

- D. Friedman et al., “In Search of the Sunk Cost Fallacy,” presented at ESA meetings Tucson 10/03, a study of behavioral biases in Web browsing.
- D. Friedman and B. Huberman, “Internet Congestion: A Laboratory Experiment,” presented at HKUST workshop 12/03 and forthcoming in a conference volume, finds (among other things) that humans can exploit “optimizing” agents by anticipating how they will react to congestion trends.
- B. Huberman et al., “Using Unsuccessful Bids to Identify Latent Demand,” Proceedings of the Computational Conference, Yale University, June 2001, a simulation study of combinatorial auctions for complex composite goods with complementarities.
- S. Anderson, D. Friedman, et al, “Buy It Now: A Hybrid Internet Market Institution,” UCSC manuscript, March 2004, an econometric study of eBay auction data.

D. Friedman, PI; Cason, co-PI: "A Laboratory Study of Customer Markets," 1997-2000 (NSF SBR96-17648, $185,000 for each PI). Friedman’s most relevant completed NSF grant compared various two-sided market formats when switching is costly for buyers. Software developed for BLN markets under this grant will be available for the currently proposed work. Results were presented at approximately two dozen professional conferences. Published papers include:

- T. Cason and D. Friedman (1999), “Customer Search and Market Power: Some Laboratory Evidence,” in M. Baye (ed.), *Advances in Applied Microeconomics*, vol. 8, (Greenwich, CT: JAI Press), pp. 71-99.
- T. Cason and D. Friedman (2002), “A Laboratory Study of Customer Markets,” *Advances in Economic Analysis & Policy* 2:1, http://www.bepress.com/bejeap/advances/vol2/iss1/art1/
- T. Cason and D. Friedman (2003), “Buyer Search and Price Dispersion: A Laboratory Study,” *Journal of Economic Theory*, 112(2), pp. 232-260.
- T. Cason, D. Friedman and G. Milam (2003), “Bargaining versus Posted Price Competition in Customer Markets,” *International Journal of Industrial Organization*, 21(2), pp. 223-251.
- T. Cason, D. Friedman and F. Wagener (2004), “The Dynamics of Price Dispersion (or Edgeworth Variations),” *Journal of Economic Dynamics and Control*, forthcoming; online since 7/04 at ScienceDirect.

E. Spohrer, PI: “East/West Consortium: Next Generation Authoring Tools and Instructional Applications,” 1994–97 (NSF CDA-9408607, $6.15 Million). As with the current proposal, this award required the creation of easy-to-use tools as well as an on-line community for sharing authoring tools and learning objects. The work helped lead to several startup next generation learning companies (CarnegieTutors, AgentSheets, Stagecast, etc.). The work lead to the Educational Object Economy (EOE.org) a non-profit which promoted and help to establish hundreds of early on-line learning communities in the 1990's, the best known today is MERLOT (Multimedia Educational Resources for Learning and On-line Teaching). MERLOT (http://www.merlot.org/) is one top on-line community that traces it lineage directly to the earlier work at Apple and then the EOE (see: http://taste.merlot.org/consortium/history.htm).

- Hundreds of publications resulted from this work -- key websites that describe these publications include: http://homepages.cwi.nl/~steven/sigchi/bulletin/1998.2/spohrer.html and http://www-jime.open.ac.uk/98/10/spohrer-98-10-06.html
- Tens of thousands of educational objects were made available to the world through this work and follow-on learning communities. Apple's on-line learning community for students and teachers

was modeled on this work, and within the first ten weeks of operation had over 10,000 registered users.

- A Google search of "Educational Object Economy" results in 3,750 hits which connect to many of the accomplishments of the "East/West Authoring Tools Group"
- The award funded one of the earliest known set of authoring tools and on-line communities for making and sharing learning objects -- the impact on future learning communities was quite substantial, including the fact that nearly a decade before the wide spread use of Wiki's the EOE website was a fully customizable website run by its users.

D.3 Proposed Research

D.3.1 Endogenous Market Entry in One-sided Hybrid Markets

As noted earlier, one-sided auctions have already been studied intensively both theoretically and empirically. However, there is an e-commerce innovation that has not yet received much attention, and it serves as an entry point to our research. A seller on eBay has the option of stating a fixed ("Buy it Now") price at which he/she is willing to sell an item before bidding begins. With Buy it Now activated, a prospective buyer either buys at the fixed price or else waits to enter the ascending-price, real-time auction. The opportunity cost of entering the auction is the difference between the value of the item to an individual and the fixed price. Buyer values are individual-specific and therefore so are buyers' opportunity costs. Individual-specific opportunity costs of entry lead to pure-strategy entry functions (Cox, Dinkin, and Swarthout, 2001), in contrast to the mixed-strategy entry functions implied by a common opportunity cost (Levin and Smith, 1994).

We will construct a Bayesian-Nash equilibrium model of the eBay ascending-price auction with the Buy it Now alternative. The model will include pure strategy entry and bid functions that predict when an item will be sold at the fixed price, when an auction will commence, and how much the bidders will bid, as determined by the probability distributions of the number of potential bidders and the item values. The model will incorporate features from models with a random number of bidders (Matthews, 1987; McAfee and McMillan, 1987b; Harstad, Kagel, and Levin, 1990) and a model with opportunity-cost-determined endogenous entry (Cox, Dinkin, and Swarthout, 2001). Theoretical comparisons will be made with the eBay auction *without* Buy it Now that will include the effects of appending Buy it Now on: (1) expected revenue to the seller; (2) market efficiency (i.e., total gains from exchange of buyers and sellers); and (3) the distribution of gains from exchange between buyers and sellers.

The theoretical models will initially assume individual private values (Vickrey, 1961; McAfee and McMillan, 1987a), in which each bidder knows the value he places on an auctioned item but knows only the probability distribution of rival bidders' values. We will later consider affiliated (roughly, positively correlated) private values (Milgrom and Weber, 1982). This sequential approach to modeling will provide insight into the effects on theoretical predictions about bidder behavior and market revenue and efficiency of the Buy it Now feature by itself and the interaction of Buy it Now with positive affiliation of item values.

The theoretical predictions comparing real-time ascending-price auctions with and without the Buy it Now alternative will be tested with laboratory experiments using the induced valuation approach briefly described in section D.2.1. The experiments will use the ABA/BAB design of auction experiments, for experimental market treatments A and B, developed in Cox, Roberson, and Smith (1982). The experiments will make possible a complete testing of theoretical predictions via induced valuation and "holding everything else constant" except for the change from (A) the ascending-price auction (by itself) to (B) the ascending-price auction plus Buy it Now hybrid market format.

Pilot research will begin on Internet auctions with treatments (A) and (B). These field experiments will be designed to be conducted on eBay by using procedures similar to Lucking-Reiley (1999) and List and Lucking-Reiley (2000). Experiments with treatment (B) involve an experimental design feature not involved in Lucking-Reiley's experiments, the choice of the specific Buy it Now price.

We will implement several fixed prices in alternative auctions of the same commodity, with the range of fixed prices based on previous eBay auction sales prices for the same commodity. A prior econometric study (Anderson, et al., 2004) of data from eBay suggests that a well-designed field experiment may produce some interesting results. For example, Anderson, et al. report that sellers with high ratings and few items to sell obtain higher prices by using Buy it Now. But the more general implications of the use of Buy it Now, including demonstration effects and learning by other sellers, effects on auction prices obtained by other sellers, and effects on entry into the market and market efficiency, need further study.

A Buy it Now, posted price alternative could be combined with a first-price sealed-bid auction, a second-price sealed-bid auction, or a decreasing-price, real-time (i.e., Dutch) auction described in section D.2.1. In the case of the Dutch auction hybrid, choice of the Buy it Now fixed price is equivalent to choice of the starting price on the Dutch auction clock, and therefore not novel. In a first-price sealed-bid auction, by contrast, it would be irrational to submit a bid higher than the Buy it Now fixed price, while the ascending-price and second-price sealed-bid auctions will, with positive probability, generate prices above the fixed price. Therefore, analysis of the properties of these hybrid formats reopens the central questions in the theory of auctions, including revenue equivalence (described in section D.2.1) or, more generally, the comparison of market formats in terms of expected revenue to the seller, market efficiency (or total gains from exchange), and distribution of the gains from exchange between buyers and sellers. Theoretical predictions comparing the eBay hybrid format with alternative hybrids will be tested with both laboratory and Internet experiments. These Internet experiments will require development of an Internet site because they will involve hybrid market formats not presently being used on commercial sites. Results from the theory and experiments with the existing eBay hybrid and the alternative hybrid market formats will be informative for both Internet firms and their customers. Thus, we will be able to compare various types of auction markets, both with and without the But it Now alternative, in terms of expected revenue, market efficiency, and distribution of gains from exchange between buyers and sellers.

The preceding extensions of market theory are motivated by recent developments in e-commerce, and are described in detail to illustrate our methods. More ambitious extensions of theory and empirics are outlined in sections D.3.2 and D.3.3, but the exposition will be more compressed.

D.3.2 Two-Sided Markets

When sellers compete for buyers, the Buy it Now feature shades into the two-sided market format called posted offer (PO). The PO allows one side to announce prices each period; the other side of the market chooses the transaction quantities. The PO predominates in traditional retail markets as well as on most B2C sites, e.g., on Half.com (recently absorbed by eBay). The PO allows buyers to enter endogenously, at their own convenience, thus saving on time costs and negotiation costs. The PO has been studied intensively in the laboratory (Davis and Holt, 1993). So far neither laboratory studies nor economic theory investigates PO markets with endogenous entry. We propose to do so by creating asynchronous trading environments that allow buyers and sellers to enter and exit at any time. We will also vary the costs for buyers to switch between markets. Entry costs are generally lower on the Internet than in bricks-and-mortar markets. Switch costs are also becoming more important on the Internet, and can be manipulated by click-throughs (lowering switch cost) and loyalty programs (raising them). Our entry studies will clarify the efficiency impact of entry and switch costs, a matter of prime theoretical and practical interest.

We will compare the efficiency of the PO to three alternative market formats when entry is endogenous. The first alternative is bilateral negotiation (BLN) or haggle. BLN requires each buyer to search for a seller (and vice versa); the pair then tries to negotiate a price and (if unsuccessful) resumes search. BLN markets are prevalent in pre-industrial retail trade, in modern B2B contracting, and in some retail Internet sites such as Priceline.com and MakeUsAnOffer.com. Laboratory research with BLN markets goes back to Chamberlin (1948), and generally shows PO markets to be more efficient than BLN. Recent work (e.g., Cason, Friedman and Milam, 2003) extends this result to markets with switch costs. It is as yet unclear whether the result holds when participation is asynchronous—indeed, the native habitat of BLN is the bazaar, which features endogenous entry and exit.

A second alternative, the continuous double auction (CDA), allows traders to make public committed offers to buy and to sell and allows traders to accept offers at any time during a trading period. Variants of CDA markets prevail in modern financial exchanges such as the New York Stock Exchange (NYSE), NASDAQ, and the Chicago Board of Trade, and are featured options on B2B Internet sites. Numerous laboratory studies beginning with Smith (1962) show that CDA markets without endogenous entry are more efficient (i.e., produce a larger total of buyer and seller trading surplus) than BLN or PO markets in a wide variety of environments. Our research will show whether the greater efficiency of CDA is robust to introduction of more realistic environments with asynchronous participation and switch costs.

The third alternative, the call market (CM), is the best understood theoretically; see e.g. Satterthwaite and Williams (1993). The CM requires participants to make simultaneous committed offers to buy or sell, and the offers are cleared once each period at a uniform price. It is used to set opening prices on the NYSE and elsewhere. Laboratory studies show that the CM is almost (but not quite) as efficient as the CDA when all buyers and sellers are present from the outset. We will compare the efficiency of the CDA to a market with equally spaced calls when buyers and sellers enter endogenously. CM could well have the advantage in this environment because the calls could attract efficient entry.

The laboratory platform for all market formats will allow participation by any combination of humans and trading agents. The agents can be programmed to make equilibrium bids and asks as in Cason and Friedman (1997), or to follow other fixed algorithms. Alternatively, they can be tuned periodically (or in real time) by humans who monitor their profitability. A popularly known example of a fixed algorithm is the "proxy bidder" for eBay's one-sided auction. The user sets only the agent's reservation price, and the agent monitors the auction. It places a slightly higher bid whenever the current highest bid is held by another bidder and is below the reservation price. As noted earlier, existing literature (e.g., Chavez and Maes, 1996; Gjerstad and Dickhaut, 1998; Cliff 2003) presents several sorts of agents for two-sided markets, and the Santa Fe Institute CDA tournament (Friedman and Rust, 1993) was a competition among many such agents of varying degrees of sophistication.

Including agents is important for several reasons. First, it provides a richer and more realistic environment for testing the relative effectiveness of alternative market formats, and it could easily overturn conventional wisdom on the subject. For example, pre-specifying agent behavior is a commitment device for the human users, as Fershtman and Judd (1987) famously demonstrated in the traditional context of Cournot competition. Cason and Friedman (1999) demonstrate that a major reason for the CDA's effectiveness is its push towards full revelation of willingness to transact, but pre-specified agents could easily resist the push. Hence the CM or the BLN format might be more effective in realistic environments.

Perhaps equally important, watching how human users tune their agents opens a unique window into human learning, strategic thinking, and motivation. Inferences regarding learning and strategic thinking come directly from user choices; by contrast, the Mouselab technology mentioned earlier requires indirect inferences from noisy information acquisition.

The point on motivation, however, is subtle and deserves a brief tangential discussion. Standard economic models assume that humans are selfish, and are unwilling to forgo own profit to help or harm others. Such models predict quite well in most two-sided markets. For example, the CDA yields prices and quantities that converge to competitive equilibrium even when virtually all of the gains accrue to one side of the market (Smith and Williams, 1989). However, models assuming purely selfish preferences fare poorly in some two-player games, e.g., the ultimatum game (Blount, 1995), the investment game (Cox, 2004), the moonlighting game (Cox, Sadiraj, and Sadiraj, 2004), and the trust game (Cox and Deck, 2003). Such data inspired the development and testing of models of distributional preferences (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Cox, Sadiraj, and Sadiraj, 2001), models of perceived intentions (Rabin, 1993; Dufwenberg and Kirchsteiger, 1999) and models incorporating both distribution and intentions (Falk and Fischbacher, 1999; Cox and Friedman, 2002). A user could program her agent to respond to helpful moves with helpful moves and respond to harmful moves with harmful moves in BLN markets and in rich variants of the CDA (already in use at the LEEPS lab) that allow traders to accept any

current bid or ask. Thus, for example, users may instruct agents to detect and defeat price discrimination schemes such as the one recently abandoned by Amazon because of consumer outrage.

Recall that the effectiveness of a market format has three aspects: static efficiency, dynamic efficiency, and evolutionary stability. The first two aspects can be assessed using classic experimental design in laboratory markets with asynchronous participation and trading agents. Tournaments will be especially helpful in assessing the third aspect. Already, agent-only tournaments have been conducted that require complete pre-specification, notably the Santa Fe Institute CDA tournament (Friedman and Rust, 1993) and the Trading Agent Competitions organized by the multi-agent community (Greenwald and Stone, 2001). The basic idea is straightforward. In each simulation run, a selected subset of agents interact in a given market environment, and the economic outcomes (especially individual profitability and market efficiency) are recorded. A balanced set of simulation runs ensures that every agent has been paired against all other agent types over a range of relevant market conditions. Thus one can identify the agent strategies that deliver consistently-high individual or market performance.

We shall design such basic tournaments for each of the online market formats. The tournaments will disclose the more profitable agent algorithms, and the resulting format effectiveness. We also will consider some more elaborate tournaments that allow human subjects to monitor profitability and to tune their agents between rounds. Beyond this are evolutionary tournaments, in which one alters the set of active agents after each round. Agents with better performance will become more prevalent, and various sorts of mutant and hybrid agents will be introduced in small proportions each round. The idea is to identify the agent strategies that survive evolutionary competition, and to evolve newer and perhaps more sophisticated agents. Finally, one can consider tournaments across market formats. One would see which formats and variants emerge from direct evolutionary competition.

Field trials for some of the more promising market formats would be a natural follow-up. Here we would seek joint corporate sponsorship. Basic designs for such trials are described in Friedman and Cassar (2004, pp. 163-4) and List and Lucking-Reiley (2000): alternative formats would be used for the identical goods in balanced sequences of trials. The field trials would test the robustness of the scientific results obtained in the laboratory and would help disseminate the findings to practitioners.

D.3.3 Markets for Services

Services are customized offerings made up of knowledge, skills, capabilities, innovations, component goods, and commodities. For example, Cisco now sells a B2B service it calls Self-defending networks, a bundle of networking hardware, software, and security management tailored for a particular customer. Wireless network connectivity is a good example of services in the B2C market. A consumer (running some particular application such as voice transmission or data streaming) has some tradeoff between price and attributes such as bandwidth, latency and jitter, and the consumer needs a connection between specific geographical locations over some specific time interval. Sellers support connectivity from geographically dispersed hubs and can estimate the cost of delivering various attribute bundles.

As we see it, the essence of the new emerging services markets is their composite and customized nature. In standard auction markets, the buyer and seller need only agree on price. In the new services markets, the buyer and seller must also agree on multidimensional product specifications. Each transaction may be for a different specification, a different bundle of attributes. How could a market function for such services?

The easiest format to imagine is a monopolist posting prices. The seller could build an algorithm to estimate marginal cost for each dimension and use a pre-established markup factor to quote an attribute-price menu to the consumer. Two problems immediately arise, but can be answered using recent research. First, what markup should the monopolist charge? The theory of nonlinear pricing (e.g., Wilson, 1997) derives profit-maximizing markups based on demand elasticities for attribute increments. In particular, Ramsey pricing identifies the constrained-efficient markups required to cover overhead as well as incremental costs. Second, as consumers constantly change location and applications, how can they best respond to changing price menus? Recent work by P. Faratin and colleagues at MIT's AI lab, e.g.,

Lee, et al. (2004), describes some adaptive learning agents that could simplify and optimize this combinatorial problem.

Absent entry barriers, a profitable monopolist will attract rival sellers, and so the PO is the first two-sided market format one might expect to emerge. But there may be more effective market formats once a sufficient number of buyers and sellers are present. Market microstructure empiricists seem to believe that as the market thickens the effective format changes from PO to CM to CDA. This is the history of the New York Stock Exchange (Schwartz, 1988), but to our knowledge the belief has never been established with any generality even in traditional markets, much less in IT-enabled markets with asynchronous participation by agents as well as humans.

The conceptual problem for service markets is that different participants will demand or supply different varieties, and so it might seem that these markets will never become thick enough for the CM and CDA formats. However, the same problem arises even for agricultural commodities, the most traditional of products, and was overcome in the 19th century by a device that we shall refer to as the *benchmark variety*. For example, a thick market evolved for hard red winter wheat #2 delivered at a particular rail junction on the seventh business day following the last trading day of March each year. Indeed, more than 50,000 futures contracts, each for 5,000 bushels, were outstanding for precisely this benchmark variety in mid-February 2004 at the Chicago Board of Trade (CBOT).[4] Most contracts are ultimately generated by producers and consumers of slightly different varieties of wheat at locations all over North America for dates several weeks earlier or later. They all trade the benchmark contract, but each wheat producer delivers his particular variety at a known price differential determined by the cost of storage and transportation. Similarly the consumers pay the price of the benchmark variety plus or minus a known price differential for the chosen variety. Most transactions do not require actual transformations into and out of the benchmark variety; active traders (the arbitrageurs) ensure that costly transformations are kept to a minimum.

Exactly the same process is at work with the Brent Crude oil contract traded on the New York Mercantile Exchange, or the 30-year Treasury Bond contract traded on the CBOT, or the new Weather Futures contracts traded on the Chicago Mercantile Exchange. The price differentials for other dates, locations and grades are known to participants, who benefit from trading in a thick market for the benchmark variety.[5]

The same pattern is emerging for new services and solutions such as computons (bundles of CPU time, storage, and system management), utility computing (bundles of consulting, software applications, and computer resources), and even "business on demand" (interoperable service bundles, e.g., for inventory and HR support, or even for manufacturing and R&D). For example, many services are sold at a benchmark rate and additional "levels" of service can be added using a Service Level Agreement (SLA). The SLA in itself provides another stage of benchmarking to simplify the complex sets of services provided (e.g., silver, gold level service). One can imagine that more detailed and complex service requirements can be priced separately and layered onto these levels.

We propose to research the question using wireless connectivity as an exemplar of the more general services problem. The new WiMax technology (802.16) can provide bandwidth of 70MB/second near its hub and significant bandwidth out to about 50 km. (By contrast, current generation WiFi is useful only out to about 50 meters.) In a basic scenario, a handful of service suppliers each have a single hub in

[4] Recent textbooks on futures markets routinely explain the CDA format used on the Chicago exchanges, and also explain some aspects of benchmark varieties and standardization, particularly the mark-to-market feature that standardizes across the dates contracts are written. However we have not yet seen a textbook or article that points out the connection, or discusses the general principle that the benchmark variety creates a thick market that can use the CM or CDA format. One of the PIs first heard the idea in a personal conversation with the late Nobel Laureate John Hicks (who alludes to it, but never spells it out, in his later books) and later encountered it as oral tradition at Chicago and UCLA.

[5] In some cases, the markets focus on the price differentials. For example, traditional bank loans to small businesses are priced as prime plus, and banks compete mainly on the "plus" differential to a specific customer rather than on the benchmark prime rate.

neighboring zip codes, and they each offer a hub price per unit of bandwidth, say for 1KB/second for one minute. More distant users would see a higher price (or equivalently, lower bandwidth or smaller units) as dictated by the physics of wireless transmission. We will induce a simple but reasonable cost structure on the service providers. On the demand side, we will induce a simple valuation structure based on the observation that some user applications (like email) are more sensitive to price and less sensitive to latency and bandwidth capacity, while others (such as videoconferencing) are the reverse. The initial market format will be PO. Users construct bundles of bandwidth over time, and purchase those that provide the best net value that they can find. We have not yet seen equilibrium derived for this basic scenario (it goes beyond Hotelling-like models in allowing variable demand quantities and inter-temporal demand substitution) but are confident that we can compute competitive equilibrium and compare it to observed outcomes in laboratory experiments.

We shall investigate several extensions of this basic scenario. First, one can introduce agents to help the buyers find the best packages. Second, one can allow sellers to offer non-linear prices, i.e., to price-discriminate by geographical location and by quantity demanded. Here again, agents tuned by sellers can give insight into sellers' motivation and strategy. As sellers compete, a benchmark variety could well emerge. That is, most sellers may choose to offer a standard size package at a particular location and to offer incremental pricing relative to that common benchmark. One hypothesis is that such practices promote closer convergence to competitive equilibrium. An alternative hypothesis is that they promote socially inefficient collusion.

Another extension we plan to investigate is alternative market formats. Buyers could also actively bid for service packages in a CM format, or negotiate with sellers in a BLN, or even enter a CDA. Our research will identify the most effective market formats in the laboratory, and will encourage pilot field trials for the most promising formats. It is worth emphasizing that reliable comparisons are almost impossible to derive from naturally-occurring data because of the strong network externality in markets: buyers prefer to enter markets with more sellers, and sellers prefer markets with more buyers. Hence whichever format appears first has a tremendous advantage. (This partly explains the persistence of 19th century market formats at the heart of New York and Chicago exchanges despite the availability of technically superior electronic technology for the last several decades.) In the laboratory one can control the market thickness and the supporting agent technology, and thus run fair comparisons of competing market formats.

A next step is to allow buyers and sellers to switch across market formats within or across trading periods. Using agents, we can control for the network externality and compare the intrinsic attraction of different market formats. This work will begin to answer the fundamental question of when more efficient market formats tend to displace less efficient formats, and when (if ever) policy interventions might help improve market efficiency. We will also examine a variety of agent interfaces (configured, attentive, and adaptive) to see which allow buyers and sellers the greatest degree of control while minimizing use-costs. Finally, we will plan evolutionary tournaments for differing kinds of agents as well as market formats.

The foregoing discussion of service markets presumes that transformation technologies exist so that a desired bundle can be obtained from a benchmark bundle at a commonly known cost. Unfortunately this is not always possible. A pervasive problem for buyers is that the value of one attribute may depend on what other attributes he can obtain. For example, the value of owning a take-off-time slot at airport A depends on which compatible landing slots can be acquired at other airports. Similar complementarities arise in valuing ownership of electrical power generators and transmission lines, or licenses for bands of the broadcast spectrum in adjacent geographical areas. The problem was first investigated in the laboratory by Rassenti, Smith, and Bulfin (1982), who developed a computer-assisted combinatorial auction for allocating airport time slots. Subsequent research includes studies of power and gas networks (McCabe, Rassenti and Smith, 1989), and railroad transportation networks (Brewer and Plott, 1996; Cox, et al., 2002).

E-commerce compounds the combinatorial problem by providing more opportunities for switching suppliers at low cost, but also potentially eases the problem by enabling automated agents to

make purchasing decisions while simultaneously solving combinatorial allocation problems.[6] Some members of our research team have recently made progress on this difficult problem. Zeng, Cox, and Dror (2004a, 2004b), develop two types of agents. One type computes optimal purchasing decisions across posted price markets with fixed transaction costs, using several efficient computational methods based on discrete location theory. The other type of agent coordinates bidding activities across multiple online auctions. The underlying coordination problem is modeled for a collection of first- or second-price sealed-bid auctions using a game-against-nature modeling approach.

The game-against-nature modeling approach taken in these agent designs ignores strategic interaction between bidding agents across markets. The full-fledged game-theoretic analysis will involve Bayesian-Nash equilibrium models. However, our initial analysis indicates that these models seem intractable for simultaneous auctions of complementary items due to non-convexities. The proposed research will explore models based on postulated "reasonable" strategies such as equal differences between stand-alone item values and bids across auction markets. Possibly, such restrictions on strategy spaces will render the problem more tractable. We will explore this question theoretically and with quantitative simulations. We also expect some insight from seeing which strategies human subjects choose in controlled laboratory experiments.

D.4 Deliverables

D.4.1 Research Papers

One series of papers will report findings on hybrid one-sided auctions with the Buy It Now (BIN) feature. One paper will derive predictions from Nash equilibrium theory for BIN-enhanced auction formats. Two or three papers will report the results of laboratory tests of those predictions for the more relevant formats, particularly English (ascending real-time), and first- and second-price sealed-bid.

A second series of papers will report findings on variants of two-sided market formats (including Posted Offer, Bilateral Negotiation, Call, and Continuous Double Auction) when participation is asynchronous and includes automated agents. The papers will adapt existing theory where feasible (e.g., for the Call market format) and will compare the formats and variants in terms of static efficiency and learning efficiency.

A third series of papers will focus on human-agent interaction in all auction and market formats. These papers will report on how human users adjust the parameters of their trading agents. As noted, key issues include strategic commitment and positive and negative reciprocity, as well as the more conventional issues of learning behavior and strategic thinking. We will begin to plan, and will seek additional funding to conduct, the field experiments and tournaments sketched in the previous section.

Work will begin on a fourth series of papers that will examine markets for customized services. Adapting agents that work well for homogeneous goods, we will induce convex preferences on human subjects for simple two-attribute goods (tentatively labeled bandwidth and latency) and let them use the agents to trade one or more benchmark varieties in Posted Offer markets. Exploratory work will commence on other market formats, more complex attribute spaces (possibly including the case of strong complementarities), and competition among potential benchmark varieties. Again, we will seek additional funding for the extensions sketched in the previous section.

D.4.2 Software Development, Infrastructure, and Dissemination

All three of the collaborating research laboratories already have some market and agent software written in the Java programming language and using XML-based messaging protocols. An early part of the effort will be to integrate all of this software into EconPort (http://www.econport.org) and thereby into the National Science Digital Library (http://www.nsdl.org). The integration should proceed relatively

[6] Most existing research is concerned with the properties of combinatorial auctions. But virtually no commerce actually occurs on combinatorial auctions. Our research will use an approach that relies on existing commercial online market formats such as posted price markets and various online auctions.

smoothly because (a) all new Java software being developed at the Economic Science Laboratory is already being integrated into EconPort; and (b) the Application Programming Interface of EconPort is designed to facilitate and guide integration of third party Java (and PHP) software. Examples of software already integrated into EconPort include: "One-Sided Auctions" for running experiments with all four of the standard one-sided auctions (discussed in sections D.2.1) with or without the software agents (or proxy bidders); "MarketLink" for running experiments with continuous double auctions (discussed in section D.2.3) and posted offer and posted bid markets; "NFG" for running normal form game experiments, i.e., multi-player strategic interdependent simultaneous move non-market decisions; "GARP" for running individual choice experiments to test the predictive power of the generalized axiom of revealed preference. An example of additional software to be integrated into EconPort by the start date for this proposal is "EFG" for running extensive form game experiments, i.e., multi-player strategic interdependent sequential move non-market decisions

Once it is stable, all new software created under the current grant also will be integrated into EconPort. Of course, EconPort will also archive all data from the research experiments. The direct benefits are clear: the software and data developed during the project will be promptly made available in the National Science Digital Library to the academic and e-commerce communities.

It should be emphasized that accessibility via EconPort will bring indirect but widespread educational and research benefits. For example, investigators with much different research questions than ours will be able to access EconPort and run new variants of hybrid one-sided auction/posted offer market formats, with or without endogenous entry and with or without activation of the integral software agents. Equally important, teachers at engineering or business schools (or in Letters and Science or policy schools or even high schools) anywhere will be able to give their students the experience of participating in continuous double auctions, call markets, and posted offer markets, asynchronously or all at the same time, and with or without software agents. Likewise, teachers, researchers and entrepreneurs interested in computons, utility computing, and business on demand will be able to peruse our data and will be able to try out for themselves original configurations of markets for new composite goods. In addition to publication of research papers and incorporation into EconPort, the PIs will disseminate research findings at professional meetings similar to those listed in section D.2.6.

.

D.5 Management Plan

D.5.1 Research Teams

The proposed research will leverage expertise and resources from three leading research laboratories and academic departments; see the facilities statements for fuller descriptions.

ESL. The Economic Science Laboratory (http://www.econlab.arizona.edu) at the University of Arizona has been a leader in experimental economics for more than two decades and currently has nine affiliated faculty members from three academic departments. In addition to its Director, ESL currently employs one Research Scientist, two Post-Doctoral Research Associates, one Support Systems Analyst, one Systems Programmer, and two graduate-student Research Assistants. Additional professional resources are available from the Artificial Intelligence Lab (http://ai.bpa.arizona.edu), the partner lab with ESL under the NSF/NSDL grant that developed EconPort. The proposed project will leverage support provided to ESL by the State of Arizona through the University of Arizona Internet Technology, Commerce, and Design Institute.

The ESL team is led by the Lab Director, PI James Cox. Other members of the ESL team have experience of recent successful collaboration with Cox, as follows. Daniel Zeng (of the Artificial Intelligence Lab) was co-PI on the NSF project that successfully developed EconPort (DUE-0226344). He is co-author with Cox on a paper studying bidding across multiple online markets for product bundles (Zeng, Cox, and Dror, forthcoming). This paper is directly relevant to the proposed project. Zeng brings expertise to the ESL team in the following areas: software agents, distributed artificial intelligence, applied optimization, and digital libraries. He will be actively involved in theory development, trading agent design, and agent implementation. Vjollca Sadiraj is co-author on several papers with Cox. She

brings expertise in both mathematical economics and experimental design to the ESL team. Vjollca Sadiraj will play an essential role in developing theory, especially in work on bidding models for environments with complementary goods, which are environments that introduce combinatorial complications into auction market environments.

LEEPS. The Learning and Experimental Economics Projects of Santa Cruz (http://leeps.ucsc.edu) Lab at the University of California, Santa Cruz has been operating steadily for 18 years, and is especially noted for experimentation with two-sided markets. It employs one senior manager/lead programmer and up to three undergraduate programmer assistants. In July 2004, the LEEPS lab, together with the entire Economics Department, moved into a new building whose other tenant is the UCSC School of Engineering (SOE). The SOE is quite new, is growing rapidly, and specializes in computer science and computer engineering.

The LEEPS team is led by the Lab Director, PI Daniel Friedman. He will be joined by UCSC Economics Professor Nirvikar Singh, who has been involved in researching and teaching about e-commerce since its earliest days. Friedman and Singh have coauthored several articles already, including two econometric studies of eBay market formats and participant behavior (Anderson et al, 2004a, b). Dr Steven Gjerstad, a leading authority on two-sided markets and recently affiliated with ESL, will join with the LEEPS lab team for its phase of the research. Associate Professor Wenjie Zhan of Huazhong University of Science and Technology, Wuhan, China, is an expert on simulating agents in CDA markets. He is visiting LEEPS lab through August 2006, and the lab anticipates future visitors interested in and able to contribute to the proposed research. Several School of Engineering faculty, including Professors Ram Akella, David Helmbold and Manfred Warmuth, have expressed interest in contributing to the design of trading agents and two-sided market platforms.

ASR. The IBM Almaden Service Research (ASR) department is located at the IBM Almaden Research Center in San Jose, California. ASR was created two years ago with the task of providing research for IBM's consulting and solutions delivery business, and to better understand the developing services economy. ASR has 37 full-time staff members including 15 PhDs with degrees in psychology, anthropology, computer science, engineering, information science, human-computer interaction, education, marketing, and operations research to name a few. ASR also has some of the best technical professionals in the IT industry, many holding masters degrees in their fields. Given the open, collaborative attitude of ASR, the broad technical and methodological capabilities of ASR members will be available to the ESL team and members of the LEEPS lab.

The Almaden team is led by PI Jim Spohrer, the director of the ASR group. The team includes three other researchers affiliated with the Almaden Research Center. Christopher Campbell has worked with Jim Spohrer for over three years on various projects and in the ASR group from its conception two years ago. He brings expertise in a number of areas directly related to the proposed project including social science methods (e.g., experimental, observational, survey, field studies), mathematical modeling, agents & agent systems, complex systems, and human-computer interaction. Ankur Chandra has been working with Jim Spohrer and Christopher Campbell for a year in the ASR group. He is currently working with Christopher Campbell on a new project to develop a large-scale agent-based infrastructure for global B2B markets and simulation research, summarized in a white paper (Campbell, Chandra, & Lehman, in prep). He has expertise directly related to this proposal; namely, expertise developing IT systems, experience creating agent technology and in web services, and a degree in economics. Nimrod Meggido works in the Computer Science theory department of the Almaden Research Center. He has a great deal of expertise related to the proposed project including game theory, reverse auctions, Dutch auctions, English auctions, fundraising auctions, sealed bid auctions, simulations, Statistical Process Control (SPC), Computability Theory.

D.5.2 Interlab Relations

The proposed work builds on ongoing collaborations among the laboratories. Cox at ESL and Friedman at LEEPS have co-authored several papers, including two in progress that also involve other ESL senior personnel. Since its inception, the UCSC School of Engineering (SOE) has cultivated close ties with units at IBM Almaden. The ASR group at Almaden has shown special interest in extending collaboration from SOE to Economics researchers at UCSC. Campbell and Spohrer have spearheaded these efforts, and by now several joint projects are under development, the current proposal being perhaps the most ambitious. Clearly there are already strong complementarities in the Economics and Computer Science expertise at the three labs. Increased synergies can be expected from close collaboration sustained over three years.

D.5.3 Coordination Plan

The ESL team will take the lead in refining software and conducting experiments for the one-sided market (or auction) formats. Researchers at Almaden (and perhaps at UCSC's SOE) will contribute to the software design, and second-site experiments will be conducted at LEEPS.

The LEEPS team will take the lead in refining software and conducting experiments for the two-sided market formats. Again, the other teams will contribute to software design, and second-site experiments will be conducted at ESL. Software and results will be disseminated through ESL's EconPort facility.

The ASR team at Almaden will take the lead in designing customizable service market software. Close collaboration with LEEPS researchers will aid the transition from simulation exercises to laboratory experiments with human subjects, to be conducted initially at LEEPS lab. When the software is ready to be installed on EconPort, second site experiments will be conducted at ESL. In parallel, ESL researchers will develop theory and software for composite goods with complementarities.

To facilitate close coordination on these projects, the three labs will sponsor a series of face-to-face meetings. For example, co-PI Zeng will spend a week at Almaden and LEEPS each year working with the local computer scientists refining agents and market platforms. Gjerstad will have extended visits at LEEPS to help develop platforms and agents for the two sided market formats. All LEEPS researchers expect to commute to Almaden (a one hour drive) at least once per month, and Almaden team members expect to increase the frequency of their visits to LEEPS.

Once projects are off the ground, daily contact across the three labs will be by email, telephone and shared files (on local servers, eventually on EconPort). As results accumulate, team members will present seminars in their home departments, at local formats (e.g., Stanford and Berkeley), at professional meetings, and at the annual HSD gatherings in Washington DC.

D.5.4 Concluding Remarks

Major social transformations of the last two centuries have largely been shaped by expanding market relations (e.g., Landes, 1998). Today is no exception, as information technology starts to transform relations among business (and other) organizations and consumers. The new market relations bring new dimensions of complexity, including hybrid market institutions with asynchronous participation by both humans and artificial agents, and transactions involving customized service bundles. To minimize unnecessary disruption, and to spread the benefits to wider segments of society, we must better understand the unfolding information economy. In this effort, existing theory can only be a partial guide. Laboratory experiments and simulations are needed to road test different market formats and agent configurations, to discover which are effective and robust, and to point the way to deeper theoretical understanding.

Our three labs propose to join forces to begin the work. The deliverables promised in section D.4 constitute fundamental research that will guide further theoretical advances, inform policy initiatives, and spur commercial development. Section D.3 also sketched further steps, beyond what is feasible given the time and funding limits of the current competition. Collaboration among our labs will continue the research after we complete the work promised here.