Experimental Business Research

Volume II

Amnon Rapoport
Rami Zwick

Springer

EXPERIMENTAL BUSINESS RESEARCH

Experimental Business Research

Economic and Managerial Perspectives

VOLUME II

Edited by

AMNON RAPOPORT

University of Arizona,
Tucson, U.S.A.
and
Hong Kong University of Science and Technology, China

and

RAMI ZWICK

Hong Kong University of Science and Technology,
China

Springer

A C.I.P. Catalogue record for this book is available from the Library of Congress.

ISBN-10 0-387-24214-7 (HB) Springer Dordrecht, Berlin, Heidelberg, New York
ISBN-10 0-387-24243-0 (e-book) Springer Dordrecht, Berlin, Heidelberg, New York
ISBN-13 978-0-387-24214-9 (HB) Springer Dordrecht, Berlin, Heidelberg, New York
ISBN-13 978-0-387-24243-9 (e-book) Springer Dordrecht, Berlin, Heidelberg, New York

Published by Springer,
P.O. Box 17, 3300 AA Dordrecht, The Netherlands.

Printed on acid-free paper

Printed in the Netherlands.

Contents

PREFACE

Amnon Rapoport
University of Arizona
and
Hong Kong University of Science and Technology

Rami Zwick
Hong Kong University of Science and Technology

This volume (and volume III) includes papers that were presented and discussed at the Second Asian Conference on Experimental Business Research held at the Hong Kong University of Science and Technology (HKUST) on December 16–19, 2003. The conference was a follow up to the first conference that was held on December 7–10, 1999, the papers of which were published in the first volume (Zwick, Rami and Amnon Rapoport (Eds.), (2002) *Experimental Business Research*. Kluwer Academic Publishers: Norwell, MA and Dordrecht, The Netherlands). The conference was organized by the Center for Experimental Business Research (cEBR) at HKUST and was chaired by Amnon Rapoport and Rami Zwick. The program committee members were Paul Brewer, Kenneth Shunyuen Chan, Soo Hong Chew, Sudipto Dasgupta, Richard Fielding, James R. Frederickson, Gilles Hilary, Ching-Chyi Lee, Siu Fai Leung, Ling Li, Francis T Lui, Sarah M Mcghee, Fang Fang Tang, Winton Au Wing Tung, and Raymond Yeung. The papers presented at the conference and a few others that were solicited especially for this volume contain original research on individual and interactive decision behavior in various branches of business research including, but not limited to, economics, marketing, management, finance, and accounting.

1. THE CENTER FOR EXPERIMENTAL BUSINESS RESEARCH

The Center for Experimental Business Research (cEBR) at HKUST was established to serve the needs of a rapidly growing number of academicians and business leaders in Hong Kong and the region sharing a common interest in experimental business research. Professor Vernon Smith, the 2002 Nobel laureate in Economics and a current member of cEBR's External Advisory Board, inaugurated the Center on September 25, 1998. Since then the Center has been recognized as the driving force behind experimental business research conducted in the Asia-Pacific region. The mission of cEBR is to promote the use of experimental methods in business research, expand experimental methodologies through research and teaching, and apply these methodologies to solve practical problems faced by firms, corporations, and governmental agencies. The Center accomplishes this mission through three agendas: research, education, and networking and outreach programs.

2. WHAT IS EXPERIMENTAL BUSINESS RESEARCH?

Experimental Business Research adopts laboratory-based experimental economics methods to study an array of business and policy issues spanning the entire business domain including accounting, economics, finance, information systems, marketing, and management and policy. "Experimental economics" is an established term that refers to the use of controlled laboratory-based procedures to test the implications of economic hypotheses and models and to discover replicable patterns of economic behavior. We coined the term "Experimental Business Research" in order to broaden the scope of "experimental economics" to encompass experimental finance, experimental accounting, and more generally the use of laboratory-based procedures to test hypotheses and models arising from research in other business related areas, including information systems, marketing, and management and policy.

Behavioral and experimental economics has had an enormous impact on the profession of economics over the past three decades. The 2002 Nobel Prize in Economics (Vernon Smith and Danny Kahneman) and the 2001 John Bates Clark Medal (Matthew Rabin) have both gone to behavioral and experimental economists. In recent years, behavioral and experimental research seminars, behavioral and experimental faculty appointments, and behavioral and experimental PhD dissertations have become common at leading US and European universities.

Experimental methods have played a critical role in the natural sciences. The last fifteen years or so have seen a growing penetration of these methods into other established academic disciplines including economics, marketing, management, accounting, and finance, as well as numerous applications of these methods in both the private and public sectors. cEBR is active in introducing these methodologies to Hong Kong and the entire Pacific Basin. We briefly describe several reasons for conducting such experiments.

First and most important is the use of experiments for designing institutions (i.e., markets) and evaluating policy proposals. For example, early experiments that studied the one-price sealed bid auction for Treasury securities in the USA helped to motivate the USA Treasury Department in the early 1970 to offer some long-term bond issues. Examples for evaluating policy proposals can be found in the area of voting systems, where different voting systems have been evaluated experimentally in terms of the proportion of misrepresentation of a voter's preferences (so-called "sophisticated voting"). In the past decade, both private industry and governmental agencies in the USA have funded studies on the incentives for off-floor trading in continuous double auction markets, alternative institutions for auctioning emissions permits, and market mechanisms for allocating airport slots and the FCC spectrum auction. More recently, Hewlett-Packard has used experimental methods to evaluate contract policy in areas from minimum advertised price to market development funds before rolling them out to its resellers, and Sears used experimental methods to develop a market for logistics.

Second, experiments are employed to test a single theory or compare competing theories. This is accomplished by comparing the behavioral regularities to the theory's predictions. Examples can be found in the auction and portfolio selection domains. Similarly, business experiments have been conducted to explore the causes of a theory's failure. Examples are to be found in the fields of bargaining, accounting, and the provision of public goods.

Third, because well-formulated theories in most sciences tend to be preceded by systematically collected observations, business experiments are used to establish empirical regularities as a basis for the construction of new theories. These empirical regularities may vary considerably from one population of agents to another, depending on a variety of independent variables including culture, socio-economic status, previous experience and expertise of the agents, and gender.

Finally, experiments are used to compare environments, using the same institution, or comparing institutions, while holding the environment constant.

3. CONTENT

Volume II contains papers under the general umbrella of economic and managerial perspectives whereas Volume III includes papers from the fields of Marketing, Accounting, and Cognitive Psychology. Volume II includes 13 chapters coauthored by 24 contributors. The authors come from many of the disciplines that correspond to the different departments in a modern business school.

In Chapter 1, Chen and Huang report on a sequence of experiments that were conducted at the Hewlett-Packard Labs, in collaboration with Ford Research Lab, to study consumer behavior in a durable goods market where leasing is prevalent. The experiments have mostly confirmed aggregate predictions of the theory and validated several qualitative features of the theoretical model. Chen and Huang observed subjects segmenting themselves into classes of behavior based on their willingness-to-pay parameters. Subjects at the low end of willingness-to-pay were priced out of both the used- and the new-goods markets. Subjects at the high end leased with increasing frequencies. They sometimes exercised their options depending on the realization of the residual quality and the potential value achievable at the used-goods market. The last segment of the subjects stayed in the middle and primarily participated in the used-goods market. The sizes of these three groups were qualitatively consistent with the theoretical predictions. Furthermore, when the strike price was increased in a different treatment, the experimental market mostly responded in the direction predicted by the model. This result is robust to small variations of market rules and sampling of subjects. Given the fact that the theoretical model has largely grossed over issues of market rules in the used-good market, the near agreement between theory and experiment is quite impressive.

On the other hand, in all the experiments the subjects with high valuation are more likely to exercise the option relative to the theoretical prediction. Chen and Huang suggest several possible explanations including risk-aversion or an ownership effects that are not addressed by the theoretical model.

Chapter 2 by Brewer begins with the question of whether it might be possible to integrate or reconcile ideas of market dynamics found in microeconomics with those found in the random walk or Martingale theory of finance.

Brewer uses a long time series generated by a Continuously Refreshed Supply and Demand (CRSD) laboratory market that provides a practical framework for an initial study of these questions. He concludes that, first, something like a random walk process can be useful in modeling the slow convergence component of prices found in CRSD markets. When a random walk in bids and asks is censored against individual budget constraints the resulting market prices appear to slowly converge towards the predictions of supply and demand. The innovative step in this model is that the random walk is not in transaction prices, but instead is a component involved in the process generating bids and asks.

The second conclusion is that the price dynamics of human-populated markets contain a number of different kinds of effects that seem to be operating simultaneously. Smoothing shows a AR(1) process similar to that seen in the constrained random walk robots. However, prices in the human-populated markets also show a complex outlier generation and correction process. A large move in prices at one trade is often corrected back towards the average with the next trade. This type of 'memory' of the process is not captured by an AR(1) statistical process or a constrained random walk of bids/asks. Removing many of the large outliers and adding an MA(1) component to absorb the remaining outlier/correction process yields an ARMA(1,1) model that varies as the market converges towards equilibrium.

A structural break in the ARMA parameters seems to occur as equilibrium is reached. The nature of this structural break is left for further research. It may suggest the use of models with multiple regimes for price discovery and equilibrium behavior rather than a simple stationary model.

Based on the above observations, Brewer speculates that a combined theory of microeconomic and financial adjustment may possibly be relied on classifying markets along several dimensions: (i) Markets with finite ending times and finite trade that can be roughly modeled as a noisy Marshallian process, and (ii) Markets with no fixed ending time and continuously refreshed supply and demand, such as the CRSD market presented in the chapter. These markets exhibit price convergence when populated by humans that can not be explained as a Marshallian process, but only as either a Walrasian class of adjustment processes or some other type of process yet to be decribed.

In Chapter 3, Saijo discusses the problem of choosing a model from several possible alternatives, and then uses global warming and emissions trading mechanism as an example. The chapter reviews three theoretical approaches to emission control: a simple microeconomic logic, a social choice concept (i.e., strategy-proofness), and a specific mechanism (Mitani mechanism) where prices and quantities are strategic variables, and the competitive equilibrium is attained via the constructed game. Since the assumed environments in the theories are quite different from one another, contradictory conclusions are derived. For example, the social choice approach presents quite a negative view of attaining efficiency, whereas the two other approaches suggest some ways to attain it. Clearly, public policy makers will be benefited from adopting a model that is a true representation of the "real" environment. However, we lack a consistent measure of proximity between the simplified theoretical assumptions and the real environment. Saijo suggests in this chapter that one way to understand (and demonstrate) how each model works is to implement its assumptions in a laboratory-based experimental environment, and he described how it has been done with the above three alternative models. The experimental approach helps drawing conclusions as to how and when theories work, conclusions that are extremely important in the public policy domain.

Chapter 4 by Friedman and Huberman reports on an experimental investigation of Internet congestion. Human players and automated players (bots) interact in real time in a congested network. A player's revenue is proportional to the number of successful "downloads" and his cost is proportional to his total waiting time. Congestion arises because waiting time is an increasing random function of the number of uncompleted download attempts by all players. The most important question in Friedman and Huberman chapter concerns rent dissipation. Would human players find some way to reduce congestion costs and move towards the social optimum, or would they perhaps create even more congestion than in Nash equilibrium? Friedman and Huberman report that human players outperform the current generation of automated players (bots). The bots do quite badly when capacity is low. Their decision rule fails to anticipate the impact of other bots and neglects the difference between observed congestion (for recently completed download attempts) and anticipated congestion (for the current download attempt). Human players are slower and less able to exploit excess capacity (including transient episodes due to random noise), but some humans are far better at anticipating and exploiting the congestion trends that the bots create. In the experiment, the second effect outweighs the first, so humans earn higher profits than bots. Overall, however, efficiency is quite low and players overdissipate potential rents, i.e., earn lower profits than in Nash equilibrium.

Friedman and Huberman conclude by offering several directions for future research including looking at "smarter" bots, connecting the results to the experiments on queuing behavior in which, contrary to the current results, fairly efficient outcomes are reported. They also propose probing the robustness of the overdissipation result by replicating the study in human-only and bots-only environments, and implementing alternative congestion functions and investigating mechanisms such as congestion taxes to see whether they enable humans and robots to earn higher profits in congestible real-time environments.

Chapter 5 presents the results of a study by Reiley on controlled experimental auctions performed in a field environment. By auctioning real goods in a preexisting, natural auction market, Reiley has collected data in a manner that is intermediate between laboratory experiments and traditional studies of field data. Given the nature of the study, some variables were unobservable and uncontrolled such as the private "valuations" of the goods. On the other hand, the procedure made it possible to hold constant most of the relevant variables in the environment, and to manipulate the treatment variable, which in this case was the existence and level of reserve prices. By giving up the ability to observe and manipulate some of variables that laboratory experimenters can control, Reiley gained a realistic environment. The participants had previous experience bidding for the types of real goods auctioned in this study (Magic cards), and the auctions took place in an Internet-based market where bidder entry decisions seemed potentially important.

Reiley reports that first, entry costs were an important feature of this real-world auction markets, thus confirming the central assumption of endogenous-entry auction theory. Second, when the same cards were auctioned twice in rapid succession, very different sets of people decided to submit bids, despite the fact that the same superset of people were invited to participate both times. This can be interpreted as evidence in favor of the stochastic (mixed strategy) entry equilibrium model, where the number of participating bidders varies unpredictably. Third, Reiley reports that, contrary to the theory of McAfee, Quan, and Vincent (1998), a zero reserve price can earn higher expected profits than a reserve price equal to the auctioneer's salvage value. Perhaps an absolute auction attracts significantly more bidder attention than an auction with even modest reserve prices, causing additional entries than might be suggested by a model of rationally calculated bidder entry decisions.

Chapter 6 by Houser and Wooders also deals with Internet auctions, investigating the effect of closing rules (hard vs. soft) on the seller's revenues. Laboratory evidence from Ariely, Ockenfels, and Roth has shown that sellers obtain more revenue when they use a soft rather than hard-close auction. This study presents evidence that the soft-close auction continues to be superior, even when it is employed in the field. Furthermore, the soft-close auction raises more revenue than a hard-close auction, even when both auctions must compete for bidders, as is the case in the field. Houser and Wooders further discuss the discrepancy between their results and the ones reported by Gupta (2001), where no difference in revenues were found between the soft and hard closing rules. They suggest that the size of the stakes may be important in understanding behavior in soft- and hard-close auctions. In particular, the revenue advantage they found for soft-close auctions may become insignificant in auctions of smaller denomination gift cards, if bidders believe that it is not worth their effort to time the placing of their bids.

In Chapter 7, Cox and Hayne investigate the conditions under which the incentive for free riding promotes rational bidding in common value auctions. Their study is motivated by the fact that, for the most part, economics has focused on models of individual rational agents whereas many important decisions are made by small groups such as families, management teams, boards of directors, central bank boards, juries, appellate courts, and committees of various types. For example, bid amounts in an economically important common value auction, the U. S. Outer Continental Shelf oil lease auction, are typically decided by committees. The Cox and Hayne approach differs from most previous research on group decision making in that they:

(a) study group decision making in the context of strategic market games, rather than non-market games against nature; and (b) use a natural quantitative measure to determine whether and, indeed, how far groups' decisions depart from rationality. Data from their previous research on group bidding behavior supports some striking conclusions. In particular, comparing bidding behavior of natural, face-to-face groups with bidding behavior by individuals reveals a "curse of information" that compounds the winner's curse. The bidding behavior of both individuals and natural groups deteriorates when they are given more information (a larger signal sample size) but bidding by groups deteriorates more dramatically. Most strikingly, natural group bidders with more information (5 signals) are significantly less rational bidders than individuals with less information (1 signal). Data from the current experiments involving cooperative and non-cooperative nominal groups reveal a rare instance in which an incentive to free ride leads to more, rather than less, rational economic outcomes. The non-cooperative nominal group treatment, with the unequal profit-sharing rule providing a free-riding incentive, produced bidding behavior that was more rational than that observed with the cooperative nominal group treatment with no incentive to free riding.

In Chapter 8, Hannan, Hoffman, and Moser discuss the comparative effectiveness of bonus versus penalty contracts. They report an experiment in which participants acted as employees under either a bonus contract or an economically equivalent penalty contract. They measured the participants' contract preference, their degree of expected disappointment about having to pay the penalty or not receiving the bonus, their perceived fairness of their contract, and their effort level. The findings can be summarized as follows: Consistent with previous work, they find that employees generally prefer bonus contracts to economically equivalent penalty contracts. However, they extend previous studies by demonstrating that employee effort is higher under a penalty contract than an economically equivalent bonus contract and that this finding is the result of two effects that work in opposite directions. The first effect is due to loss aversion, which makes employees more averse to having to pay a penalty than not receiving a bonus, and causes them to choose more effort under the penalty contract. The second effect reflects reciprocity, which causes employees who consider their contracts to be fairer to choose more effort. Because employees generally perceived the bonus contract to be fairer than the penalty contract, reciprocity caused employees to choose more effort under the bonus contract. They find support for both of these opposing effects, with reciprocity dampening, but not completely offsetting, the dominant effect of loss aversion on employee effort.

Hannan, Hoffman and Moser conclude by discussing the implications of their results for explaining why in practice most actual contracts are bonus contracts rather than penalty contracts. In particular, they point out that their results show that conventional economic analysis fails to capture either employees' preferences for bonus contracts or the fact that penalty contracts motivate higher effort. They speculate that there are still other costs associated with offering penalty contracts, or benefits associated with offering bonus contracts, that are not yet reflected in any of the currently available explanations for why penalty contracts are rarely observed in practice. For example, in some work settings, employees have ways to retaliate against firms that offer a penalty contract other than to withhold effort. In such settings, employees could work hard to avoid the penalty, but then quit the firm to work for another firm; or, alternatively, they could withhold effort and possibly pay the penalty, but then extract monetary benefits from the firm through other means (e.g., employee theft) to make up for their lost incentive compensation. If firms anticipate such retaliation, they may conclude it is cost effective to offer employees bonus contracts rather than penalty contracts.

Chapter 9 by Croson and Schinnar reports on a study that experimentally tests the impact of managerial incentives on competitive (market) outcomes. They use a symmetric Cournot duopoly setting with perfect information and no uncertainty and compare different compensation schemes; one in which managers are paid as a function of the profits of the firm, and a second where they are compensated based on their performance relative to the other firm in their industry. When managers are compensated based on firm profits, the equilibrium of the game involves collusion. However, when managers are compensated based on relative profits, the equilibrium devolves to the perfectly competitive outcome. They test this simple theory in an experiment.

The experimental results support the model's comparative-static predictions: how managers are compensated (based on absolute or relative profits) has important implications for collusive behavior. In addition to validating the theory, these results have important lessons for antitrust regulators. To determine whether an industry is collusive it is not sufficient (and may not even be necessary) to look at the industry's output; one should also look at managerial incentives of the individual firms. Similarly, regulating managerial incentives may have a bigger impact than simply denying specific mergers. Even in very concentrated (two-party) industries such as the one implemented in the current research, when incentives were relative rather than absolute, outcomes were competitive. Thus, even in industries where concentration and other usual measures of collusive potential are the same, the amount of inefficiency that is observed is likely to depend on the incentives of the managers.

Chen in Chapter 10 studies the dynamic stability of Nash-efficient public goods mechanisms and reconciles theory with previously reported experimental results. Until now, Nash implementation theory has mainly focused on establishing static properties of the equilibria. However, experimental evidence suggests that the fundamental question concerning any actual implementation of a specific mechanism is whether decentralized dynamic learning processes will actually converge to one of the equilibria promised by theory. Based on its attractive theoretical properties and the supporting evidence for these properties in the experimental literature, Chen focuses on supermodularity as a robust stability criterion for Nash-efficient public goods mechanisms with a unique Nash equilibrium. Her paper demonstrates that given a quasilinear utility function the Groves-Ledyard mechanism is a supermodular game if and only if the punishment parameter is above a certain threshold value while none of the Hurwicz, Walker and Kim mechanisms is a supermodular game.

The Falkinger mechanism can be converted into a supermodular game in a quadratic environment if the subsidy coefficient is at least one. These results generalize a previous convergence result on the Groves-Ledyard mechanism, and are consistent with the experimental findings. Two aspects of the convergence and stability analysis in this paper warrant attention. First, supermodularity is sufficient but not necessary for convergence to hold. It is possible that a mechanism could fail supermodularity but still behaves well on a class of adjustment dynamics, such as the Kim mechanism. Secondly, The stability analysis in this paper, like other theoretical studies of the dynamic stability of Nash mechanisms, have mostly been restricted to quasilinear utility functions. Consequently, the maximal domain of stable environments remains an open question. Results in this paper suggest a new research agenda that systematically investigates the role of supermodularity in learning and convergence to Nash equilibrium.

In Chapter 11, Bearden, Rapoport, and Seale study entry times in queues with endogenous arrivals, and in particular the dynamics of play on the individual and aggregate levels. In previous studies the authors (and several additional co-authors) have investigated experimentally how delay-averse subjects, who patronize the same service facility and choose their arrival times simultaneously from a discrete set of time intervals, seek service. Taking into account the actions of others, whose number is assumed to be commonly known, each self-interested subject attempts to maximize her net utility by arriving with as few other subjects as possible. Each player can also stay out of the queue on any particular trial. Using a repeated game design and several variants of the queueing game, the authors report consistent patterns of behavior (arrival times and staying out decisions) that are accounted for successfully by the symmetric mixed-strategy equilibria for the games, substantial individual differences in behavior, and learning trends across iterations of the stage game. The major purpose of the chapter is to account for the main results of several different conditions by the same reinforcement-based learning model formulated at the individual level.

The authors adopt a “bottom-up” approach to explain the dynamics of the repeated interaction. The focus is on the distributions of arrival time on both the aggregate and individual levels. They begin the analysis with a simple model that has as few parameters as possible, and modify it in light of the discrepancies between theoretical and observed results. The performance of the learning model is mixed. It accounts quite well for the aggregate distributions of arrival time in four of the five conditions and produces heterogeneous patterns of individual arrival times that are quite consistent with those produced by the experimental subjects. However, the learning model generates considerably more switches in arrival times than observed in the data and somewhat smaller mean switch magnitude than observed in all the experimental conditions.

Chapter 12 by Schotter surveys a number of papers all of which have investigated the impact of advice on decision-making. In general, this advice is offered by decision makers who are only slightly more experienced in the task at hand than are the people they advise. Such an advice is referred to in the chapter as “naïve” advice. Despite this lack of expertise Schotter reports a number of common findings. First, people tend to follow the advice offered to them. This is seen in a number of ways. For example, in Ultimatum and Trust games the amounts of money sent by the senders to the receivers is remarkably close to the amounts these subjects are advised to send. In coordination games, subjects tend to choose the action they are told to even when that action differs from the action that constitutes a best response to their beliefs about their opponent. Second, not only is advice listened to and followed, but it tends to change people's behavior. Games played with advice are played differently than those same games played without it. In addition, efficiency is generally higher when games are played with advice. This is true in coordination games where subjects tend to coordinate more often, in social learning tasks where advice increases the incidence of herds and cascades but always on the right decision, in one-person learning tasks, and finally in the Minimum Games when advice is public. Schotter proposes that the reason why advice is so beneficial is that it forces decision makers to look at the problem they are facing in a more detached manner. The act of giving advice forces one to rethink the problem at hand while the act of receiving advice forces one to evaluate the advice that is given. Both endeavors lead a decision maker to take a more global approach to the problem and help a decision maker see the forest rather than the trees.

The last Chapter (Chapter 13) by Lee, Weg, and Zwick reports on the failure of Bayesian updating in repeated bilateral bargaining game. They study a game that allows for reputation building. Of course, it is quite natural for people or institutions to misrepresent their true nature in pursuit of gaining some benefits which otherwise could not be attained. Although misrepresentation may touch on questions of the law there are situations in which misrepresentation may only be a matter of benign convenience and opportunity as the framework explored in the present chapter shows. The basic setting for this study includes a buyer and a seller. The seller possesses five units of a product that he intends to sell to the buyer in five periods, one unit in each period. The buyer is known to the seller to be one of two types: low cost (L) and high cost (H) with probabilities π and $1 - \pi$, respectively. This is operationalized as the low or high costs related to the seeking of an alternative supplier for an identical product that the seller proposes to sell. Upon receipt of the proposal to sell the product at a specific price, the buyer may accept it and thus terminate the transaction, or opt to search (at a cost) for a better price by another supplier. The search for another supplier is always successful; however, the price may be better or worse than the current one proposed by the present seller. If the buyer elects to search, she abandons the opportunity to purchase the unit at the original seller asking price and is committed to pay the "searched" price even if it is higher than the current asking price (i.e., this is a no recall environment). The game is repeated (5 times) among the same two players. The equilibrium of the game is very similar to that of the game described by Kreps and Wilson and also to the one that was experimentally tested by Camerer and Weigelt. Whereas Camerer and Weigelt concluded that "sequential equilibrium describes actual behavior well enough" the experiment reported in this chapter demonstrates the limit of the above conclusion.

In particular, the ultimatum nature of the basic game tends to overwhelm rational behavior on the part of the sellers, and buyers are not cognizant of favorable prices occurring later in the game. The authors conclude by discussing the sources of difficulties in playing the game "correctly" and offer several suggestions for further theoretical developments to accommodate the current findings.

Similar to the first volume, volumes II and III should be viewed as work in progress and guide for future research. The conference and the resulting book were designed to provide a place for intellectual exchange of ideas between experimentalists within the various business disciplines. We hope that the exposure we have provided for the experimental method in business will inspire the reader to pursue the method and take it to new heights.

ACKNOWLEDGEMENTS

We owe thanks to many for the successful completion of volumes II and III. Most importantly, we express our gratitude to the contributors who attended the conference and participated in insightful discussions. The conference was supported financially by a grant from the Hong Kong University Grant Commission to cEBR (Project No. HKUST-3, Experiential based teaching for networked economy), and by an RGC Direct Allocation Grant (Project No. DAG02/03.BM78) to Rami Zwick and Soo Hong Chew. Additional financial support was provided by HKUST. Special thanks are due to Professor K C Chan the Dean of the HKUST Business. We wish to thank Maya Rosenblatt and Maggie Chan, the conference secretaries, without their help the conference would have been a total chaos, and Chi Hang Chark for the splendid and dedicated work in preparing and formatting all the chapters for publication. We also thank Kluwer for supporting this project.

Chapter 1

DURABLE GOODS LEASE CONTRACTS AND USED-GOODS MARKET BEHAVIOR: AN EXPERIMENTAL STUDY

Kay-Yut Chen
Hewlett-Packard Laboratories

Suzhou Huang
Ford Motor Company

Abstract

Leasing has become an increasingly prominent way for consumers to acquire durable goods such as automobiles. How markets respond to changes in lease contracts has enormous implications to producers such as Ford Motor Company. In this paper, an experimental model was developed to study the interaction between lease contracts that embed an option to purchase and an underlying used-goods market. Experiments with subjects playing roles of heterogeneous consumers have confirmed many salient features predicted by the theoretical model. These features include the segmentation of subjects into classes of behavior, and directional response to pricing in the used-good market to the provision in lease contracts.

1. INTRODUCTION

Leasing has become an increasingly prominent way for consumers to acquire durable goods. Very often, lease contracts embed options that allow lessees the right but not the obligation to purchase the item at the end of the lease. This form of lease contract is very popular in the automobile industry. In this paper, an experimental model was developed to study how this kind of lease contracts interacts with an underlying used-goods market. This research, although self-contained, is the first stage of collaboration between HP Labs and the Ford Motor company to create a general framework to address some of the unique issues in automobile marketing.

The standard option pricing theory approach assumes perfect competition and frictionless market (Black and Scholes 1973, Merton 1973). In this framework, agents are assumed to be homogenous and non-strategic, and transaction costs are all negligibly small. Furthermore, producers are assumed to have little market power

A. Rapoport and R. Zwick (eds.), Experimental Business Research, Vol. II, 1–19.

and hence are treated as price takers. These assumptions are quite reasonable for financial and well-traded commodity markets. However, they often fail to capture the key features of those markets in which durable goods are leased such as automobiles and heavy machineries. In these durable-goods markets, consumers are heterogeneous and need to act strategically in their consumption decisions in a time-consistent manner due to sizable transaction costs that they have to incur for trading used goods. Similarly, many of these durable goods are often made by a few big producers with a high degree of differentiation. This, in turn, implies that, rather than simply acting as price takers, the producers can enjoy certain market power in pricing their goods and changing provisions in lease contracts. Aiming to gain insights, Huang and Yang (2002) had constructed a theoretical model that explicitly incorporates many of the salient features of the economic environment that is more appropriate for these durable-goods markets.

However, some issues remain unresolved if the purpose is to adapt the insights to make policy decisions in the real world. One key issue is whether the model is robust with respect to the stringent rationality requirements imposed upon the consumers. Huang and Yang (2002) employed the solution concept of the Markov perfect equilibrium (Maskin and Tirole 1988). The solution requires the consumer to have perfect knowledge of the present and future prices, as well as the supply and demand of used goods, which are all endogenously determined by solving complicated mathematical equations. This is obviously beyond the undertaking of an average consumer. Even if every agent in the system can perform the mathematics required, there is ample evidence to show that people are neither risk-neutral nor even adhering to expected utility maximization (Camerer 1995). Another question is whether the theoretical results are robust with respect to variations of the price discovery process, which can vary depending on the particular market mechanisms used. It is also impractical trying to infer the answer from real world data because there are many unobserved or uncontrollable variables. The most promising approach is laboratory experiment.

A series of experiments was conducted at HP Experimental Economics Lab. In each experiment, around 23–28 subjects were recruited to play the role of consumers in a hypothetic durable-goods market. Standard experimental economics procedures were followed while there was a slight variation to the standard design of treatments. We gave subjects exact information about the experiment. They were told that their monetary rewards depended on their aggregate performance of the experiment. We preserved anonymity with respect to roles and payment and we used no deception. The experimental model was directly adapted from the setting in Huang and Yang (2002). There was one brand of homogenous goods when they are new. Each unit of goods was associated with a quality measure. A new unit always started with the quality of one. In the first period, a random amount was consumed. The residual or leftover quality was observed at the end of first period, which later would be consumed in the second period. We chose this particular structure to capture the characteristics of automobile market. Given a brand, new cars are generally identical. Thus, all new units started with the same quality that is normalized to one. The primary measurement to determine the relative value of a used car of the same brand is its usage such as mileage driven. In any given period of time, the actual mileage accrued is uncertain at the inset of a lease contract to the lessee. This is why we chose that the quality consumed in the first period is uncertain ex ante. On the other hand, when a lessee is deciding whether to exercise the option at the lease end, the residual quality is ex post and hence is treated observable to the lessee. For simplicity, we assume that each good has a lifetime of two periods. This implies that any residual quality is completely consumed in the remaining life span of the used good, i.e., the next period after its lease.

Each subject is only allowed to have at most one unit of the good, new or used. A positive value is given to the subject at each period if he or she owns a unit. This value is a function of the quality consumed and a private parameter called the willingness-to-pay. Each subject had a different willingness-to-pay parameter, which was chosen to span uniformly over an interval. To focus on issues related to lease we limit the producer only lease its products, and hence outright selling is left out of the scope of this study. In each period, a subject had four alternatives: start a new lease, purchase a unit from the used market, exercise the option to purchase the leased unit at the end of the term if the subject was a lessee in the preceding period, or hold no unit. If a subject started a new lease, he would consume a random amount (the mean and variance of this amount were common knowledge) in the same period. In the end of the period, he would face the choice of whether to exercise his option to buy the used unit with the strike price that was specified by the lease contract. If not, the unit would be returned to the producer and sold in the subsequent used-goods market.

The new-good market is modeled as a fixed take-it-or-leave-it lease contract. The lease term is one period. The lease price and the strike price were common knowledge and remained constant throughout each experiment. Since almost all automobile lessors have been using auction (to dealers, not consumers) as the standard method to re-market used cars, we have decided to use a round-based ascending bid auction for the used-good market. All the supply in the used-good market came from returned off-lease units. Thus, both the size and prices of the market were endogenously determined. We assume that the residual quality of a good in the used-goods market is observable in our experiment, and thus we sidestepped the adverse selection problem that was made notorious by lemons in used-car business between individual sellers and buyers. The rationale behind our choice is based on the following two facts. First, the most relevant parts of the used-car market for auto producers are those related to off-leases or fleet rentals that are relatively new, typically one to two years old. Second, before the auction process, almost all used-cars are inspected and the results are well documented and disseminated to any potential buyers; and any deals that are in dispute can be conveniently settled through arbitrations.

The experimental observations have largely confirmed the qualitative features of the theory, both at the aggregate and individual levels. Furthermore, the comparative statics of the experimental market in responding to the change of the strike price are consistent with that of the theoretical model. On the other hand, the experimental results also provide evidence suggesting that there are systematic biases in the theoretical model due to the assumptions of perfect rationality and risk aversion.

The remainder of the paper is organized as follows. In section 2, we first recapitulate the theoretical model introduced by Huang and Yang (2002), and then briefly outline how the model is solved. In section 3, we spell out the details of the experimental design. The experimental results along with comparison with theoretical predictions are presented in section 4. We conclude and point out some future directions in section 5. The appendix provides some additional results that are relevant for the discussions in the main text.

2. THE THEORETICAL MODEL

The content of this section is extracted from Huang and Yang (2002). All details can be found in the original paper.

2.1. *The Goods and Lease Contract*

All the goods have a lifetime of two periods. They are regarded as homogeneous when they are new. This allows us to normalize the quality measure for the entire life span of the goods to be 1. Depending on the usage of a good in its first period, the residual quality of the good in the second period is denoted by $\delta \in (0, 1)$. Since the usage of a particular good is uncertain at the onset of the lease contract, δ is treated as stochastic for lessees[1] and is assumed to obey an exogenous distribution with a known density $g(\delta)$.

For convenience, we take this distribution as a lognormal distribution with the following mean and volatility parameters: $\mu = -1$ and $\sigma = 0.2$. We further define $\phi(x) = \int_0^x g(\delta)\, d\delta$ and $\Phi(x) = \int_0^x \phi(\delta)\, d\delta$. On the other hand, when a used good enters the used-good market, δ is treated as observable for all participants there. We further assume that any remaining quality of a used good is completely consumed in the second period.

The lease contract allows the lessee to use a new good for one period with a lease price of r. At the lease end, the lessee has the option to either keep the used good by paying a pre-determined strike price k or returns the used unit to the producer without additional obligation.

2.2. *Consumer Preference*

From the consumers' point of view, the goods are differentiated vertically. Consumer's heterogeneity is parameterized by θ, representing the willingness to pay for a unit of quality. We assume that the distribution of θ is uniform on [0, 1] and does not change over time. In the context of the experiment, each individual will have the same θ for the whole experiment. The only uncertainty an individual can encounter is when he leases a new good: he is unsure of the residual quality δ of the leased good at the lease end.

Table 1. The utility flow matrix with a state s and an action a for consumer θ: $\Pi_\theta[s, a]$

$s \setminus a$	L_δ (with unknown δ)	C (with known δ')	U_δ (with known δ)
$L_{\delta'}$	$(1-\delta)\theta - r$	$\delta'\theta - k$	$\delta\theta - q(\delta)$
$\bar{L}$	$(1-\delta)\theta - r$	$-\infty$	$\delta\theta - q(\delta)$

The utility flow of an individual at each period is a function of δ and θ, as well as a function of pertinent prices: lease price r, strike price k, and used-good price $q(\delta)$ for a unit of used good with residual quality δ. To simplify the setting, we assume that the transaction cost for a consumer to sell used goods is prohibitively high. We further exclude the outright selling of new goods, in order to focus on studying the optionality embedded in the lease contract. No outright selling also implies that there is no trade-in. The following table details the assignment of the relevant utility flows for consumer θ in the case when all new goods are only leased. Since we are dealing with durable goods with transaction cost, the utility flow will have to be explicitly state dependent.

In the above table, $L_{\delta'}$ denotes a state that the consumer leased a new unit in the last period and has a known residual quality δ' entering the current period. $\bar{L}$ represents any state that the consumer did not lease a new good in the last period. L_δ (with unknown δ) depicts the action of leasing a new good in the current period with a consumed quality of $1 - \delta$. C (with known δ') signifies the action of exercising the option to keep the used good of residual quality δ' that was leased in the last period. U_δ (with known δ) is the action of buying a used good with an observed residual quality δ. The lease price r and strike price k are announced by the producer at the beginning of every period, and are kept constant throughout the experiment.

One can easily recognize that consumers are assumed to be risk neutral in the theoretical model. While simplifying the mathematical treatment, some of the detailed quantitative discrepancy between the theory prediction and experimental observation may be attributed to the risk neutrality assumption.

2.3. Consumer Behavior: Theory

We will only be concerned with the steady limit of the dynamic equilibrium where player's reaction function becomes independent of time, and each consumer adopts a constant consumption pattern (a fixed sequence of strategies).

Concept of Solution: The dynamic aspect of the consumers' decision-making is modeled using the solution concept of Markov perfect equilibrium developed by Maskin and Tirole (1988). Strategies that a consumer can take depend only on the current state. A general equilibrium is embedded into the game at every period to endogenize the used-goods market in a style of Huang, Yang and Anderson (2000). The used-goods price is determined by the clearance condition for each realizable residual quality in the used-goods market. Grossing over the microeconomic process of price formation is again for the technical tractability. Therefore, we should expect the theory to make sense only on average, and anticipate that experimental results, which are obtained from explicitly treating the used-goods market as an ascending auction, are going to deviate from theory predictions at some detailed level. For the justification of how a competitive price emerges from auction processes, readers are referred to Wilson (1977) and Milgrom (1981).

Consumers' Bellman Equation: Given the various prices in the steady limit, consumer θ at state s solves the following Bellman equation

$$V_\theta[s] = \max\{E_\delta[\Pi_\theta[s, L_\delta] + \rho V_\theta[L_\delta]], \Pi_\theta[s, C] + \rho V_\theta[C], \max_{\delta \in \Delta_U} (\Pi_\theta[s, U_\delta] + \rho V_\theta[U_\delta])\},$$

where $\rho \in [0, 1]$ is the discount factor and Δ_U stands for the set of realizable residual qualities in the used-goods market. The first term in the curly brackets corresponds to leasing a new good, the second term corresponds to exercising the option, and the third term corresponds to buying a used good. That an expectation with respect to δ appears only when the consumer chooses to lease reflects the fact that the residual quality δ is ex ante for new leases and is ex post for actions associated with used goods.

Consumer segmentation: When the exogenous parameters of the model are appropriately chosen, the above Bellman equation admits a unique solution with the following consumer behavior. Consumers are naturally segmented by a pair of division points θ_m and θ_M (with $0 < \theta_m < \theta_M < 1$). Low valuation consumers, $\theta \in (0, \theta_m)$, choose to stay out of the market. Consumers in (θ_m, θ_M) choose to buy used goods. High valuation consumers, $\theta \in (\theta_M, 1)$, choose to lease new goods or exercising options according to the reaction function

$$R_\theta[L_\delta] = \begin{cases} L, & \text{if } \delta < \zeta(\theta) \\ C, & \text{if } \delta > \zeta(\theta) \end{cases}.$$

This threshold rule leads to the following probabilities for consumer θ to be in leasing a new good or continuing to consume the used good by exercising the option: $h_L(\theta) = 1/[2 - \phi(\zeta(\theta))]$ and $h_C(\theta) = [1 + \phi(\zeta(\theta))]/[2 - \phi(\zeta(\theta))]$. The values of the division points are determined from the conditions that consumer θ_m is indifferent in staying out of the market or buying a used good with an arbitrarily low residual quality, and that consumer θ_M is indifferent in buying the used good with the highest realizable residual quality in the used-goods market δ_M or leasing a new good.

Average Payoff Function: The average payoff function per period $(1-\rho)V_\theta[.] \equiv \tilde{V}_\theta[.]$ can be shown to have a finite limit when $\rho \to 1$:

$$\tilde{V}_\theta[I] = 0 \qquad \text{for } \theta \in (0, \theta_m);$$

$$\tilde{V}_\theta[U_\delta] = \delta\theta - q(\delta) \qquad \text{for } \theta \in (\theta_m, \theta_M);$$

and

$$\tilde{V}_\theta[L_\delta] = \tilde{V}_\theta[C] = \frac{[1 + \Phi(\zeta(\theta)) - \zeta(\theta)\phi(\zeta(\theta))]\theta - [1 - \phi(\zeta(\theta))]k - r}{2 - \phi(\zeta(\theta))}, \quad \theta \in (\theta_M, 1).$$

As it is well-known, the $\rho \to 1$ limit is called time-average criterion. We will justify later why this criterion is the relevant one for the experimental setting.

Option-exercising Threshold: The option exercising threshold is related to the average payoff function as $\zeta(\theta) = (k + \tilde{V}_\theta[C])/\theta$. In the limit of $\rho \to 1$ the threshold satisfies the simple equation:

$$r - k = [1 + \Phi(\zeta(\theta)) - 2\zeta(\theta)]\theta.$$

Used-good Market: The used-goods supply is from off-leases. The price for a used good with residual quality δ is determined by the clearance condition:

$$\theta(\delta) - \theta_m = \int_{\theta_M}^{1} d\theta \frac{\phi(\min\{\delta, \zeta(\theta)\})}{2 - \phi(\zeta(\theta))}.$$

Supplementing the clearance condition with the equation of marginal substitution rate $\theta(\delta) = \dfrac{dq(\delta)}{d\delta}$ and the terminal condition $q(0) = 0$, the price for a used good with residual quality $\delta \in (0, \delta_M)$ can be written as

$$q(\delta) = \theta_m\delta + \int_0^{\delta} d\delta' \int_{\theta_M}^{1} d\theta \frac{\phi(\min\{\delta', \zeta(\theta)\})}{2 - \phi(\zeta(\theta))}.$$

Residual Quality Distribution in Used-goods Market: The residual quality distribution in the used-goods market is modified from the original residual quality distribution according to

$$\tilde{g}(\delta) = g(\delta) \int_{\theta_M}^{1} d\theta \frac{I[\delta < \zeta(\theta)]}{2 - \phi(\zeta(\theta))} \Bigg/ \int_{\theta_M}^{1} d\theta \frac{1}{2 - \phi(\zeta(\theta))},$$

where $I[\cdot]$ is the indicator function. The return rate is obtained by integrating over δ,

$$\omega(r, k) = \int_0^1 d\delta \tilde{g}(\delta) = \int_{\theta_M}^1 d\theta \frac{\phi(\zeta(\theta))}{2 - \phi(\zeta(\theta))} \Bigg/ \int_{\theta_M}^1 d\theta \frac{1}{2 - \phi(\zeta(\theta))}.$$

Numerical Solution: Some of the equations, such as the clearance condition and threshold equation, do not appear to be amenable in closed form. However, they can be easily solved numerically.

3. EXPERIMENTAL DESIGN

The experimental model was implemented in the HP Experimental Economics Software. Every experiment has around twenty five subjects each playing the role of a consumer who procures a durable good that lives for two periods. Each person is limited to process at most one unit of good in any given period. Instructions for the experiments were posted on the web. Each subject had to pass a web-based quiz before he was allowed to participate. The instructions and the accompanying quiz are available at: http://www.hpl.hp.com/econexperiment/lease/instructions.htm

3.1. Preference

Preferences were induced according to the vertical differentiation model described above.

The homogeneity of the new goods means that we can normalize the quality measure for the entire life span to be 1. The residual quality for a used unit is denoted by $\delta \in (0, 1)$. This parameter divides the whole quality into new and used. δ is drawn from a known distribution.

Consumer's heterogeneity is parameterized by $\theta \in [0, 1]$, which represents the willingness to pay for a unit of quality. The theoretical analysis assumes that θ is drawn from a uniform distribution and that θ does not change over time. The experimental design deviates slightly from these assumptions. We chose θs to span over the [0, 1] interval in the following manner. Consider an experiment with N subjects. The interval [0, 1] was divided into N equal intervals: $[0, 1/N]$, $(1/N, 2/N]$, . . . $((N - 1)/N, N]$. Each of these "mini" intervals was then assigned to a different subject. Each subject's θ was drawn randomly from his "mini" interval. This design ensured we would observe θs to span uniformly over the [0, 1] interval.

Each subject was allowed to switch θ once in each experiment. This was done usually on period 13. If a subject drew his first θ from the interval $(m/N, (m + 1)/N]$, his second θ would be drawn from the interval $((N - m - 1)/N, (N - m)/N]$. Thus, if a subject had a very low first θ, his second θ would be guaranteed to be high. This was done to address fairness concerns.

Furthermore, each subject was given \$1 for each period he completed since only a subset of the subjects were expected to make money.

3.2. Decisions

In each period, an individual has the following choices of action: 1) hold no unit; 2) lease a new unit; 3) purchase the leased unit with the strike price k if at a lease-end; and 4) buy a used unit from auctions held by the producer.

New units were leased from an exogenous source, referred to as the producer in the experiment. The lease price and the strike price were announced to the subjects in the beginning of each period and they understood that these prices were common knowledge. In all the experiments, both the least price and lease-end strike price stayed the same throughout the experiment. The lease price was paid at the onset of the lease contract and entitles the lessee to use a new good for one period. At the beginning of the 2nd period of the life of the leased good, the individual has two alternatives. He could purchase this unit at the strike price or he could decide to return the unit to the producer. The residual quality is unknown at the time of signing the lease contract and becomes known at the time of deciding whether to exercise the option. If a unit was returned to the producer, it would be sold in an auction in the following period.

Simultaneous round-based ascending bid auction was chosen as the auction mechanism. This is similar to the actual used car auctions that most of the auto producers are conducting, except that in our case the units are directly sold to the consumers while large scale used-car auctions are only open to dealers. Since auction was also well studied in the literature (Kagel et al. 1995), this choice also enables us to interpret our results in the context of past experiments if the need arises. Residual qualities were announced before each unit was auctioned. Furthermore, subjects did not have to decide on a new lease or exercising of an option until the end of the auction.

3.3. Treatments

The key issue of importance to the producer's used-car remarketing business is how the endogenous used-car market reacts to changes in the lease contracts. In particular, the producer is interested in the effect of the lease-end strike price. Two treatments were used in the experiments, one with a strike price equal to $k = 0.08$ and the other with a strike price of $k = 0.16$.

4. RESULTS

4.1. Overview

As outlined in the Introduction, there are dual motivations for carrying out the experimental study. The first is to check whether the economic assumptions adopted by the theoretical model, such as perfect rationality and risk neutrality, are plausible. The second is to gauge the robustness of the theoretical predictions on the price formation mechanisms in the used-good market given that not all the assumptions of the model would hold true when real human beings are involved. Due to its lack of

Table 2. Summary of the four experiments

Experiment	*Number of subjects*	*Number of periods*	*Lease price*	*Strike price*	*Number of new leases*	*Number of auction units*
1	25	25	300	80	209	54
2	28	24	300	160	158	85
3	28	23	300	160	174	122
4	23	21	300	160	151	95

an explicit modeling of the microeconomic process in the used-good market, we can only hope that the theoretical model makes quantitative sense on average at best. In addition, the number of subjects is about 25 for each experiment, which may not appear to be small from the first sight. However, taking into account the consumer's heterogeneity and the complexity of possible consumption decisions that can endogenously emerge; we still expect substantial finite-sample fluctuations. Therefore, when contrasting the experimental results with their theoretical counterparts, we will mostly concentrate on qualitative and comparative static aspects.

For the sake of convenience, all valuations and prices in the experiment are rescaled by a factor of 1000, so that subjects can submit bids that are integers. A total of 4 experiments were conducted. The following table provides an overview of all the experiments.

Ideally, more experiments would be conducted. However, business constraints only allowed for 4 experiments in this study. Despite the small number, experimental results seem to be robust with respect to some qualitative features of the model.

Two issues arose in the course of this work that resulted in the choice of parameters (three experiments with strike price of 160). The first issue is sampling effects. Typical experiments at HP Labs use Stanford students who already had prior experience through participating in other earlier experimental economics projects. Around the same time when Experiment 2 was conducted, HP had expanded its scope of subject recruitment to a local city college, mainly due to the need of a completely different project. As a result, some students from the local city college were conveniently recruited for Experiment 2. These students from the local city college never had any prior experience. Unfortunately, this fact was overlooked during the process of training subjects to become proficient in decision-making in the context of this experiment. As a consequence, significant portion of the subjects earned substantially less than predicted, whereas this was not true for the rest of the subjects in the same experiment or in other experiments (see Appendix). These subjects had a strong inclination to participate and to win auctions even when their best option was to choose other strategies.

A second issue was a good illustration of how the detail in market institution design matters. The rule of the game in the first three experiments was that a player could choose either to bid in an auction *or* to start a new lease but not both. In effect, at the last round of auction, a player has to choose between bidding, which could result in no win, and start a new lease. Consequently, some high valuation consumers, who decided to participate in the auction process but were not able to win the bid in the final round, were deprived of the chance to lease new goods. Obviously, the theory has no bearing on this issue since it is based on a market clearing and does not specify an explicit mechanism of how the clearance is achieved. This issue turns out to be not severe for Experiment 1, because the size of the used-good market is small, but it is quite noticeable in Experiments 2 and 3.

These two issues compounded together resulted in fewer new leases in Experiments 2 and 3, in which some of the high valuation subjects achieving very low payoffs. This prompted us to change the game rule in Experiment 4 from "either to bid or to lease new" to that a subject can always have a chance to lease new if he loses in the auction. We believe that the latter rule is closer to reality. Thus, we conducted 3 experiments using $k = 160$ with slightly different market rules and different samples of subjects. Ideally, we would conduct more experiments to contrast the effects of these issues. However, the rigorous time table of a business related project did not allow us to do that. Furthermore, while this is not standard experimental methodology, we do not believe these variations substantially altered our conclusions.

Before we present the details, it is important to emphasize that there are no free parameters in the theoretical predictions when they are compared with their corresponding experimental counterparts, once the exogenous parameters, such as the distributions of consumer's heterogeneity and residual qualities, are chosen to be the same. However, finite sampling sizes in the experiment can introduce systematic bias to these distributions. To partially alleviate this kind of bias, especially for the distribution of residual qualities, the theoretical predictions are calculated based on the values of parameters computed directly using the finite samples realized in the experiments. Due to our way of sampling θ, finite sampling effect for consumer heterogeneity is less of a problem. Finally, we fix the value of the only behavior parameter in the theory as $\rho \to 1$, or according to the time-average criterion. This choice can be justified by the fact that subjects' monetary rewards are mostly based on their cumulative performance in more than 20 periods. Thus, the experimental setting is such that subjects are motivated to maximize their average payoff per period.

4.2. Aggregate Level Comparison

Among the aggregate variables that we examine are the following: 1) new-lease probability per period per consumer, which serves as a measure for the demand of new goods; 2) return rate, which measures how likely a lessee exercises the embedded option in the lease contract; 3) average used-good price, which is endogenously

Table 3. Theoretical predictions with the intended parameters of the residual quality distribution: $\mu = -1$ and $\sigma = 0.2$

Theory	*New-lease prob. per period per consumer*	*Return rate per lease*	*Average used-good price*	*Aggregate surplus per period per consumer*	*Producer revenue per period per consumer*
$k = 80$	0.34	0.53	105	104	135
$k = 160$	0.33	0.89	126	96	144

determined; 4) aggregate surplus per period, which measures how consumers as a whole benefit from participating in the market; and 5) producer revenue per period, whose sources of contribution include new-leases, exercised options and resale of used goods. Variables normalized by the number of periods and/or number of subjects will enable us to combine results obtained from different experiments of the same setting, and to compare results from different experimental settings in a meaningful manner.

In order to have an appreciation of how finite sampling correction affects the theoretical prediction, we first list these predictions with the originally chosen parameters for the residual quality distribution $\mu = -1$ and $\sigma = 0.2$ in Table 3. Typically, the finite sampling implies about 5% corrections to the mean and 10% corrections to the volatility. As we will see shortly, all aggregate variables, except return rate, are not very sensitive to the finite sampling correction.

Table 4 lists the results of Experiments 1 to 4, along with the corresponding theoretical predictions corrected by the finite-sampling effect. Since Experiments 2, 3 and 4 share the same $k = 160$, we first average the aggregate results from these three experiments and then compare the average to the theory. The differences between these three experiments also serve as a crude measure of behavior fluctuations from rather small sample sizes of subjects. Given the fact that there is no fitting process involved in the comparison, the level of the agreement between experimental results and theoretical predictions in Table 4 is quite remarkable. Quantitatively, the worst case is the return rate, in which the experimental values are systematically lower than that of the theory by about 30%. One way to interpret this systematic difference is risk aversion. The only uncertainty in this model is the consumption in the first period of a new lease, represented by an unknown residual quality that is only realized at the lease-end. Thus, risk averse agents may be inclined to keep the leased unit, whose value is known at the time of exercising the option, instead of starting another new lease. Consequently, return rate will be lower than the theory that assumes risk neutral consumers. Another possible way to interpret the systematic discrepancy may be traced to ownership effects. However, to settle the true cause, additional theoretical modeling and experimental investigation are needed.

Table 4. Experimental results and theoretical predictions with the finite-sample parameters of the residual quality distribution realized in each experiment

Experiment	*New-lease prob. per period per consumer*	*Return rate per lease*	*Average used-good price*	*Aggregate surplus per period per consumer*	*Producer surplus per period per consumer*
1 (k = 80)	**0.33**	**0.26**	**115**	**94**	**130**
Theory[2]	*0.33*	*0.37*	*113*	*101*	*130*
2 (k = 160)	0.24	0.54	147	52	107
3 (k = 160)	0.27	0.70	122	70	117
4 (k = 160)	0.31	0.63	90	88	130
Average (2, 3, 4) (k = 160)	**0.27**	**0.62**	**120**	**70**	**118**
Theory[3]	*0.32*	*0.80*	*132*	*91*	*142*

A primary policy question that a producer is interested in is how the market would respond to a change in the strike price. The theory predicts that an increase in the strike price from $k = 80$ to $k = 160$ at a fixed lease price will lead to a slight decrease in total lease volume, a substantial increase in the return rate, an increase in average used-good price, a reduced aggregate surplus for consumers, and an increase in producer revenue. All these directional changes are confirmed in Table 4, with the exception of producer revenue, which went the opposite way of the theoretical prediction. We attribute this deviation to the fact that there are too few new leases in Experiments 2 and 3, caused by issues of market rules and subject sampling mentioned earlier. It is worth noting that the theory predicted a substantial change only in the return rate while all other changes are more moderate. Experimental results confirmed this substantial change in the return rate.

We chose not to report standard deviation statistics. Since the game is dynamic in nature, data across periods were not independent. Thus, calculating standard deviations, or any other variance estimates, across periods would not be useful. Furthermore, variations in subject behavior were mostly driven by their different willingness-to-pay parameter θ. Therefore, reporting variance estimates across individuals would not truly reveal heterogeneous individual characteristics such as risk aversion. However, most of the comparative static holds true between any of Experiment 2, 3, or 4 (with $k = 160$) and Experiment 1 (with $k = 80$). Thus, we have some confidence that the comparison is valid.

4.3. Detailed Level Comparison

We now examine how the experimental results and theoretical predictions compare at a detailed level. In particular, we are interested in seeing how patterns of consumer behavior emerge as a function of willingness-to-pay. We are also interested in seeing how used-good prices change with variations of residual quality. For the sake of space limitation, we will only use the results for Experiment 1 as illustrating examples. In most of the cases, the results of Experiment 1 are quite typical. Due to the fact that the used-good market is treated tersely in the theory, we expect that the theory will fare less well at a detailed level than at an aggregate level.

In the following we treat the same subjects with a different θ essentially as a different consumer. If all the data were used, each subject would yield two points. Thus, we observe a total of twice as many consumers as the number of subjects in each experiment. It can be argued that the data in the first two periods with freshly assigned θ values should be thrown away because of start-game effects. However, we found that the conclusions are not dependent on whether we exercise this option.

4.3.1. Average Payoff and Used-good Price

Figure 1 shows average payoff per period as a function of consumer heterogeneity θ. In the left panel of the figure, the theoretical payoff curve tracks very closely the experimental payoffs. The right panel of the figure indicates that the observed used-good prices are clustered around the theoretical prediction. The trend that higher residual quality implies a higher used-good price is reproduced, though with large fluctuations. There is a small number of observations whose residual qualities are higher than the point where the theory curve ends. This signals a slight behavior deviation from the theory, which predicts that there is an upper limit in residual qualities in the used-good market due to the presence of the option. Nevertheless, Figure 1 allows us to conclude safely the following results.

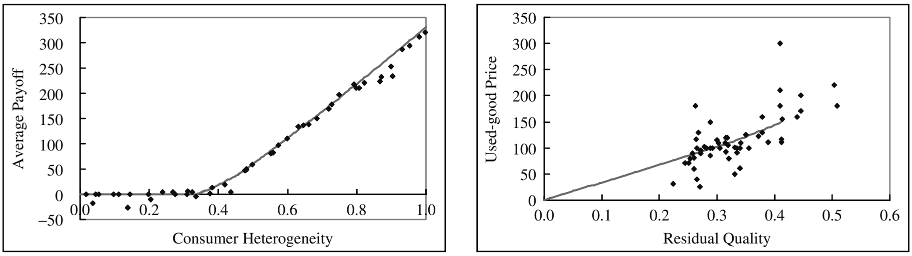

Figure 1. Average payoff as a function of consumer heterogeneity (left panel) and used-good price as a function of residual quality (right panel). Curves are theoretical predictions and diamond points are experimental observations in Experiment 1.

Result 1: Observed payoffs are consistent with the theory.

Result 2: Observed used-good prices are consistent with the theory.

4.3.2. Behavioral Segmentation

The theoretical model predicts that subjects would be segmented endogenously into three classes of behavior. Lower valuation consumers $\theta \in (0, \theta_m)$ are priced out of the market. Medium valuation consumers in $\theta \in (\theta_m, \theta_M)$ participate in the used-good market. High valuation consumers $\theta \in (\theta_M, 1)$ lease new goods and occasionally exercise the option embedded in the lease contract at lease-end.

Behavior segmentation can be captured in two measures: new-lease probability and auction-winning probability. Figure 2 shows these probabilities as functions of θ. In Experiment 1, the theory predicts $\theta_m = 0.33$ and $\theta_M = 0.47$, respectively. As one can see from Figure 2, both new-lease probabilities and auction winning probabilities are quite low when $\theta < 0.3$. This supports the conclusion that on average, low valuation consumers are priced out of the market. New lease probabilities begin to rise at around $\theta = 0.4$ and become quite close to the theoretical curve from around $\theta = 0.5$ onward. On the other hand, though still roughly concentrating at around the right region, auction-winning probabilities are much more spread than the theory's prediction. From time to time, consumers who would be theoretically the pure used-good buyers also enter the new-lease market, and consumers who would be theoretically pure lessees venture into the used market. One interpretation is that the fundamental economics forces were operating correctly. However, the perfect rationality assumption in the theory is obviously violated, leading to the smearing in consumer segmentation.

Interestingly, the smeared behavior does not cause a substantial payoff gap, as can be inferred from the left panel in Figure 1. This implies that the economic incentive that is responsible for the sharp segmentation in theory is not very strong

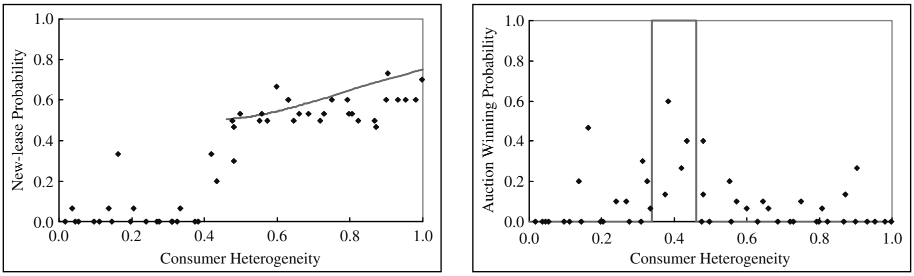

Figure 2. New-lease probability (left panel) and auction winning probability (right panel) as functions of consumer's heterogeneity. Lines are theoretical predictions and diamond points are experimental observations in Experiment 1.

for those consumers whose willingness-to-pays are in the middle, and occasional "mistakes" are gracefully tolerated. In addition, Figure 2 also provides evidence on why several subjects have their payoffs much lower than the theoretical curve. For example, consumers whose θ values lie between 0.8 and 0.9 should have leased more new goods rather than participated in auctions. Nevertheless, the following conclusion can be drawn.

Result 3: Strong but Imperfect Patterns of Behavioral Segmentation.

4.3.3. Cherry Picking

Theoretically, units with a higher residual quality have a higher chance of being purchased by the consumer exercising his lease-end option. Thus, the units returned to the producer would have a distribution skewed towards the low-end compared to the original distribution of residual qualities. This kind of *cherry picking* phenomenon is also observed in the experiment. Figure 3 shows the distribution of residual qualities for all the units and the distribution for those units that were returned to the producer and subsequently entered the used-good market. Notice that not all high residual quality units were returned to the producer as predicted.

Furthermore, Kolmogorov-Smirnov Tests (Table 5) show that, in three out of four experiments, the distribution of residual qualities of the returned units is consistent with model predictions. Experimental evidence not only confirms the cherry picking phenomenon in a *qualitative* fashion, but also suggests that the theory is sound *quantitatively* despite all the handicapping factors mentioned before.

Result 4: Cherry Picking Observed and Consistent with Theory.

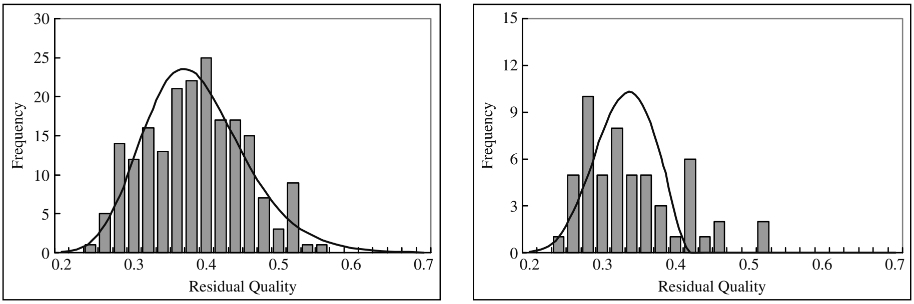

Figure 3. Distributions of residual qualities for all used units (left panel) and for those that enter the used-good market (right panel). Bars are experimental observations in Experiment 1, and curves are theoretical predictions, which are normalized to have the same masses as in the experiment.

Table 5. Kolmogorov-Smirnov Test to see if residual qualities of the returned units were consistent with the theoretical distributions

Experiment	*Observations*	*K-S Statistics*	*P-Value*
1	54	0.177	0.97
2	85	0.092	0.78*
3	122	0.069	0.70*
4	95	0.057	0.48*

* cannot reject the null hypothesis at 95% confidence that observed residual qualities obey the distribution specified by the theoretical model.

5. CONCLUSION

A sequence of experiments was conducted at Hewlett-Packard Labs, in collaboration with Ford Research Lab, to study consumer behavior in a durable goods market where leasing is prevalent. The experiments have mostly confirmed aggregate predictions of the theory and validated several qualitative features of the theoretical model. We observed subjects segmenting themselves into classes of behavior based on their willingness-to-pay parameters. Subjects at the low end of willingness-to-pay were priced out of both the used- and the new-goods markets. Subjects at the high end leased with increasing frequencies. They sometimes exercised their options depending on the realization of the residual quality and the potential value achievable at the used-goods market. The last segment of the subjects lived in the middle and primarily participated in the used-goods market. The sizes of these three groups were qualitatively consistent with the theoretical model. Furthermore, when we increased the strike price in a different treatment, the experimental market mostly responded in the direction predicted by the model. This result is robust even with small variations of market rules and sampling of subjects. Given the fact that the theoretical model has largely grossed over issues of market rules in the used-good market, the near agreement between the theory and experiment is highly non-trivial.

On the other hand, in all the experiments, the subjects with high valuation are more likely to exercise the option relative to the theoretical prediction. There are multiple possible explanations. One such possibility is risk aversion that is not addressed by the theoretical model. With risk aversion, a leasing subject has the tendency to keep the used unit that entails no uncertainty relative to lease a new good that has an unknown consumption in the first period. Other explanations such as ownership effects may also account for the discrepancy between theory and experimental results. More evidence is needed to pinpoint the correct explanation.

The effect of learning in the experiment appears to manifest mostly in whether subjects are used to the economic context of the experiment. Once subjects familiarize themselves with the decision-making process, there is no obviously discernable effect associated with progressive stages of the experiment. However, due to the complex setting of the experiment, less experienced subjects, as exemplified in Experiment 2, took a long time to figure out what they ought to behave and hence earned significantly less payoffs comparing to more experienced subjects.

There are several directions that can be viewed as natural extensions of the current work. To settle whether the aforementioned systematic bias in return rate is caused by risk aversion or something else can be pursued by extending the theory to include risk aversion and conducting additional experiments that are specifically designed for this purpose. Another interesting direction is to treat the residual quality being only partially observable, which in turn will allow the possibility of studying the interplay between optionality and adverse selection. Investigations of lease contracts with more sophisticated options and under oligopoly market structure are other topics for future exploration. In addition, it is important to realize that the setting of the current experiment is not very far from many realistic business environments. Adapting the experiment described in this paper to field studies has the potential to provide useful business insights. Finally, work has already begun to use a modified version of this experiment to examine business strategies in other aspects of the automotive market.

NOTES

[1] If the residual quality were known to the lessee at the signing of the lease contract, there would have been no risk factor in each consumer's decision-making process, at least theoretically. This, in turn, would have made the option embedded in the lease contract meaningless.

[2] The finite-sample parameters of the residual quality distribution realized in Experiment 1 are $\mu = -0.95$ and $\sigma = 0.18$.

[3] The finite-sample parameters of the residual quality distribution realized in Experiments 2, 3 and 4 are $\mu = -0.96$ and $\sigma = 0.22$.

REFERENCES

Black, F. and Scholes, M., (1973). "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81, 637–659.

Camerer, C., "Individual Decision Making." In *The Handbook of Experimental Economics*, edited by Kagel, J. and Roth, A. Princeton Univ. Press, 1995.

Huang, S. and Yang, Y., (2002). "Pricing Lease Contracts with Options in Imperfect Markets of Durable Goods." Technical Report, Ford Research Laboratory.

Huang, S., Yang, Y. and Anderson, K., (2001). "A Theory of Finitely Durable-Goods Monopoly with Used-Goods Market and Transaction Costs." *Management Science*, 56, 549–569.

Kagel, J., "Auctions: A Survey of Experimental Research." In *The Handbook of Experimental Economics*, edited by Kagel, J. and Roth, A. Princeton Univ. Press 1995.

Maskin, E. and Tirole, J., (1988). "A Theory of Dynamic Oligopoly: I and II." *Econometrica*, 56, 549–569 and 571–599.

Merton, R., (1973). "Theory of Rational Option Pricing." *Bell Journal of Economics*, 4, 141–183.

Milgrom, P., (1981). "Rational Expectations, Information Acquisition, and Competitive Bidding." *Econometrica*, 49, 921–943.

Wilson, R., (1977). "A Bidding Model of Perfect Competition." *Review of Economic Studies*, 44, 511–518.

APPENDIX: PAYOFF CURVES

The following figures show payoffs of all four experiments. Each point represents the average payoff of a subject under the same willingness-to-pay parameter. It is interesting to note that earnings for all subjects in Experiments 1 and 4 and for most subjects in the other two experiments are very close to the predicted values. As pointed out in section 4, some subjects in Experiments 2 and 3 were earning substantially less money.

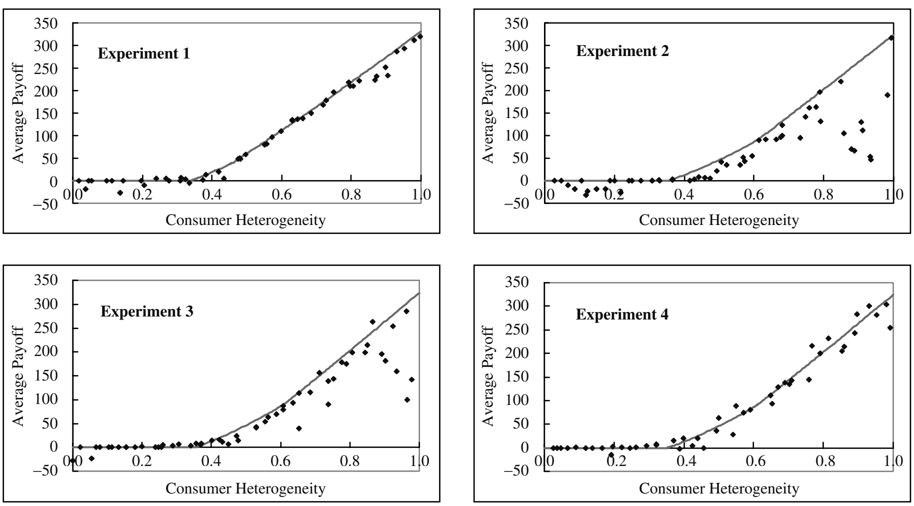

Figure 4. Average payoff as a function of consumer heterogeneity in all four experiments.

Chapter 2

TOWARDS A HYBRID MODEL OF MICROECONOMIC AND FINANCIAL PRICE ADJUSTMENT PROCESSES: THE CASE OF A MARKET WITH CONTINUOUSLY REFRESHED SUPPLY AND DEMAND

Paul J Brewer
Hong Kong University of Science and Technology

Abstract

Microeconomics and financial economics provide alternative models of market dynamics. A long history of laboratory results shows that market prices in the laboratory converge towards the static predictions of microeconomic theory with a resulting classical efficiency of allocation. Yet, the informational efficiency of market prices, often treated as a starting axiom for financial market theory, requires instead that current prices represent fair gambles over an unknown distribution of future prices: financial price processes are idealized as random walks with independent increments perhaps modified by some notion of heteroskedasticity such as stochastic volatility. Unlike prices following a Marshallian path, random walks do not generally converge towards an equilibrium price. The conflict between these two views of market processes is explored and a model that is a hybrid of the microeconomic and financial approaches is constructed and compared against data from laboratory markets involving continuously refreshed supply and demand.

1. INTRODUCTION

This chapter is a very rough first attempt to integrate some ideas from across microeconomics and finance about the price dynamics of competitive markets. The research is from the point of view of an experimental economist interested in laboratory market equilibration, not from the point of view of general asset pricing or finance in general. The goal is not to resolve all the questions one might have about the nature of price dynamics, convergence or the differing approaches or assumptions that may be involved across various fields.

A. Rapoport and R. Zwick (eds.), Experimental Business Research, Vol. II, 21–45.

Instead, the goal is more modest, to put forward the notion that the noisy equilibration of a fairly simple single market is still a subject worthy of study. There are no "states of the world" in the sense of classical finance and, correspondingly, no laboratory bets on securities whose values are based on coin tosses or dice rolls. Instead, there is a pair of markets, a private market and a public market. Buyers and sellers receive private, seemingly random opportunities to buy or sell a good from the "experimenter" in their private market and are able to trade with each other in the public market. Subjects are not told anything about the distribution of these opportunities. The supply and demand curves representing the aggregate of these private market opportunities are held stationary and the experimenter observes the time series of voluntary trading prices in the public market.

Since the market participants do not know ex-ante what the public market price should be, there is a kind of endogenous heterogeneity and complexity of beliefs and knowledge about market conditions more typical to the experimental economics literature than the classical finance literature. It is the general success of experimental economics in providing a means of studying this peculiar kind of complexity that is hoped to make such a laboratory approach worthwhile.

Although a broad view of some of the problems one encounters in merging ideas from different fields is important, ultimately the research reported here is much more narrowly focused upon a particular data set and a particular form of time-series analysis. One can then attempt to ask questions about the adequacy of simple, stationary models: Can price equilibration be described by a simple mathematical equation with fixed parameters or is a model with two or more regimes more appropriate? Does something happen when markets equilibrate that we can detect in the time-series properties of the data? The data reported here is an attempt to get at these questions, among others. The research is not expected to answer many questions at this stage, but instead it is an attempt to stimulate new questions and to begin a long process of obtaining answers made possible through the continued work of future researchers.

The remainder of this section will provide an overview of some literature, but does not pretend to be a guide to this subject for newcomers nor can it even hope to even briefly credit all those whose research formed the present understanding of markets. The introduction concludes with a brief road map organizing the research to be presented.

Early laboratory studies into market behavior, beginning with Smith (1962), were not designed merely to confirm or demonstrate known principles of economics. Early experimental environments by design violated three common assumptions once thought appropriate for the applicability of competitive models: (i) *perfect information* was violated as student subjects typically knew only their own costs and values when trading, not the costs or values of others, the aggregate supply and demand, or the distributions from which costs and values were drawn; (ii) *continuity* was violated at the unit level and the agent level because the units traded in the market are indivisible and because agents were not trading small quantities relative to the aggregate market; (iii) *perfect rationality* was probably violated because the student subjects often had no previous exposure to trading and therefore could not be expected to trade as well as the perfectly rational *homo economicus*, and possibly not even as well as a businessman or professional trader.

High efficiencies of allocation and convergence of observed prices and traded quantities to the predictions of competitive theory nevertheless occurred in early laboratory markets. A critic who only saw the experiments as a misconstruction relative to the requirements of existing theories, inadequately "simulating" a larger market, or relying "too much" on data from students instead of experienced businessmen or professional traders would potentially miss an interesting result regarding the robustness of competitive processes. These early laboratory markets were real markets. The results showed that the details of trading institutions matter: the prices in the markets of Chamberlin (1948) did not converge nearly as well as Smith (1962) because Smith included specific kinds of trading structure – the publicly observable bids and asks recorded on a blackboard and the improvement requirement (bids go up, asks come down until a trade occurs) inherent to double auction rules – while Chamberlin's completely unstructured approach left traders on their own to decide what to do as they walked around and searched out trades with others in the room. Charles Plott and a number of other researchers duplicated Smith's early laboratory results, going on to laboratory investigations involving multiple markets, transformation and production, and other complex scenarios.

The broad pattern of these results demonstrated that markets could function fairly well – given the proper structure and a bit of learning by repetition – with lumpy goods and only a few inexperienced traders in a variety of situations and applications. Plott (2000) has argued that laboratory research on market processes and equilibration can support a modernization of Hayek's view of the market. Hayek (1945) viewed markets as human institutions providing a means of imperfect, but self-correcting, coordination and solution to a demand/supply problem without having to convey all the information about market conditions to a single mind. Previous laboratory studies reveal market equilibration likened to a rational but almost mechanical process, possibly unrecognized by the market participants, attempting to find the solution of an equation balancing supply and demand. Even though no one (except the experimenters) knows the equations or has full knowledge of the parameter values needed to solve the equations, the rationality inherent in profit-seeking behavior would drive the process to equilibrium.

In contrast, the framework Gode and Sunder (1993) developed as an alternative explanation for market equilibration by way of their Zero Intelligence (ZI) robot algorithm demonstrated a strong potential for a mechanical, non-rational convergence processes based only on budget constraints and not on profit maximization. The ZI robot framework is still a popular environment for beginning a study of more complex phenomena (see Farmer, Patelli and Zovko (2004) for a recent example or Duffy (2004) for a review). Prices in markets populated by the ZI robots appear to converge towards competitive equilibria and exhibit negative autocorrelation of price changes. Cason and Friedman (1996) find negative autocorrelations of price changes in laboratory markets populated by inexperienced human traders, with the autocorrelations moving towards zero and positive autocorrelation with more experienced subject pools. Both studies also show that large surplus trades would occur earlier in the market, with convergence being driven by the fact that this leaves the price-constrained low surplus trades to occur later in the market period.

While laboratory microeconomics has developed a body of empirical regularities surrounding imperfect – but functional – markets, standard financial theory generally begins with a set of axiomatically defined perfect markets and derives further properties under various conditions in an uncertain world using mathematical probability theory. Take, for example, the case of a popular and regularly traded stock on the NYSE or NASDAQ. Analysts follow the business closely. Therefore, at least among the major market participants setting prices, one might assume perfect information or at least homogeneous information. Millions of shares are traded, so continuity is virtually satisfied. The major market participants are generally expert traders, and so should be acting rationally. If one assumes that all known information has been fully processed by a perfect market, the prediction of finance in the short term is amazingly simple: the share price should represent a fair gamble based on the probability distribution of possible share prices in the near future.

Over time, prices should exhibit the properties of a Martingale process, such as zero autocorrelation of price changes. Field tests on financial market data yield various non-zero results.[1] However, a careful theorist can still argue that the Martingale property is an *ex-ante* property related to *historical expectations* about future prices and therefore impossible to test *ex-post* based solely on observed prices alone without some additional assumptions – see for instance, Bossaerts (2002; pp. 42–43). Without certain simplifying assumptions, one would would instead need to be able to somehow record what the "market was thinking" about future prices, and test whether the price at each moment in time equaled this expectation as beliefs evolve.

Of course, this ex-ante kind of Martingale theory is much more difficult to falsify, and also causes the fine details of beliefs to become important. Are beliefs actually homogeneous so that all market participants have the same expectations or is this merely a convenient approximation? If beliefs are homogeneous, are they correct or at least unbiased? Does it matter if homogeneous, correct beliefs do not initially exist but do form over time as the market converges? Or do the correct beliefs exist because the market participants exist in a world of stationary probabilities where the frequency of various kinds of events, and their effects on prices, are well known? Without evoking criticism of any microeconomic or financial theory and shying away for now from the technical details that make the approaches of microeconomics and finance to equilibration so different, it is interesting to note that the Hayekian view of market equilibration as a process of solving for prices without conveying all the necessary information to a single mind is in such stark contrast to the view of more widely studied theories in finance that assume that all market participants are indeed of a single mind in the sense of holding identical, correct beliefs. These questions are already well known but are tricky and quite technical to deal with, and are beyond the immediate scope of this work.

Several earlier laboratory experimental approaches to financial economics are reviewed by Sunder (1995). Information aggregation from insiders to the general market and belief formation were common areas for exploration. More recently, Bossaerts (2002) reviews many theoretical issues in finance and discusses laboratory experiments structured specifically to test asset-pricing models in a multi-asset risky environment. Bossaerts (2002; p. 129) notes that laboratory markets converge slowly, and this slow convergence in prices may require models with adjustments or biases of "market" beliefs away from the perfect beliefs assumed by the efficient market hypothesis.

Most previous work in finance-based laboratory experiments, including the work cited above, required experiments with many markets and many uncertain states of the world in order to fit the mold of the financial models. Instead, the research to be reported here focuses on equilibration of a single market. The connection to finance is in the efficient market hypothesis and its implication for Martingale or random walks in prices.

Irregardless of whether changes in financial market prices are due to random shocks to the profitability of an underlying business or random noise traders, if there is a pattern to the price changes then there is a potential for profit that should not exist in a perfect market. The Martingale or random walk hypothesis can be thought of as an axiomatic description of perfect market prices without reference to an underlying firm or asset or any specific requirement limiting the scope to only financial markets.

Does the notion of price as a Martingale process apply to laboratory markets? If prices do not follow a Martingale or random walk, is the notion of a random walk still useful somehow? Can the random walk somehow be reconciled with the notion of an imperfect market that is attaining competitive equilibrium over time? The answer to the first question will be no, both on principle and empirically prices in laboratory markets clearly do not follow a Martingale process. But the initial answer to the latter two questions will surprisingly be yes.

The process of reaching this result is as follows: Section 2 describes the Continuously Refreshed Supply and Demand (CRSD) Environment that is used to generate long data sets and disrupt the means by which ZI robot populated markets converge to equilibrium. Human-populated CRSD markets still appear to converge towards an equilibrium price, so something more is happening with the humans that do not happen with the ZI robots. Section 3 identifies the microeconomic and financial approaches to market convergence. Section 4 compares and contrasts these two approaches and identifies some issues that would appear to prevent the financial model from describing the behavior of laboratory markets. Section 5 shows how to use the random walk to design a new kind of trading robot that captures some of, but not all of, the dynamics of the human-populated market in the CRSD environment. Markets populated by the random walk robots show price dynamics that can be described fairly well as an AR(1) process. However, markets populated by humans also show a kind of outlier-correction whereby prices deviate from the convergence path and then pop back up to near the previous price. Outliers and corrections can be modeled as a type of MA process so that the joint process becomes an ARMA process. Section 6 analyzes ARMA models of the price convergence of human-populated markets and summarizes the findings. Section 7 discusses conclusions.

2. THE CONTINUOUSLY REFRESHED SUPPLY AND DEMAND ENVIRONMENT

Figure 1 shows a set of instantaneous supply and demand curves that are held constant in the continuously refreshed supply and demand (CRSD) experimental study of Brewer, Huang, Nelson and Plott (2002). The environment is implemented by means of a set of java-based programs accessed from a standard web browser such as Microsoft Internet Explorer. Human traders sitting at a web browser see their screen divided into a public market, for trading within the group, and a private market, which displays a set of private trading opportunities (production costs or redemption values for a single unit of good) available only to that subject. Subjects complete trades and make money by arbitraging their private market prices available from the experimenter against the public market prices available from interaction with other experimental subjects. For example, if a subject can buy a

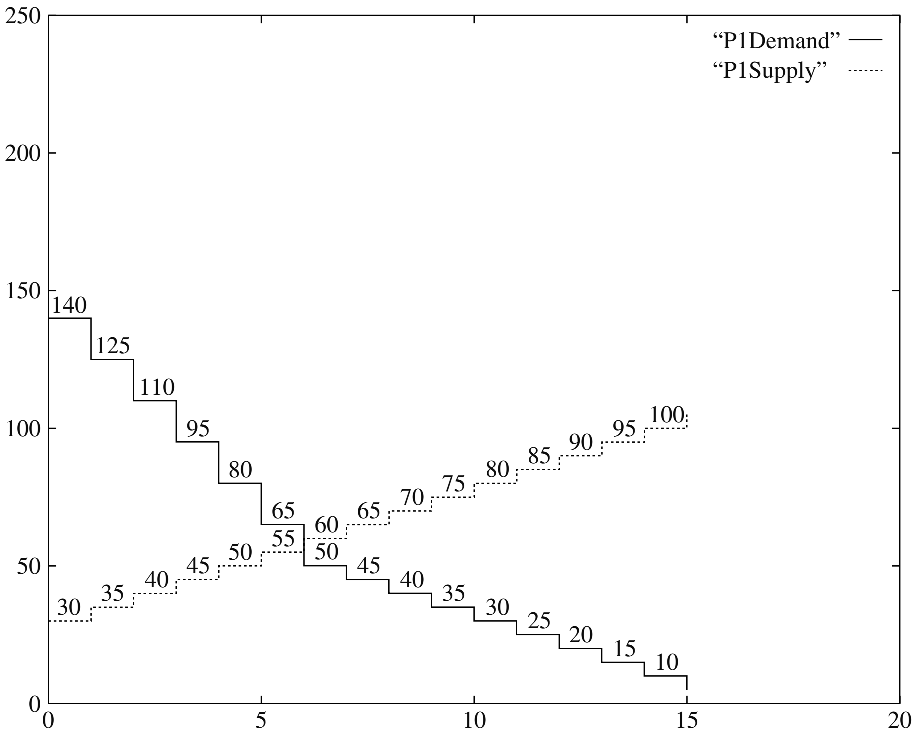

Figure 1. Sample Supply/Demand Environment.

unit in the public market for \$50 and sell it in the private market for \$70, they will earn a profit of \$70 − \$50 = \$20. Net profits are paid in cash at the end of the experiment.

The costs and values in Figure 1 are distributed among the subjects via the private markets visible on their individual trading screens and are recycled among the subjects as trade occurs. The details of this recycling will be explained below. The experimenter is not primarily concerned with this private market recycling, but instead the focus is on observing the trading between the humans in the public market.

Continuously refreshed supply and demand is a technique for recycling the costs and redemption values in a double auction experiment. In contrast to standard double auction experiments where gains from trading are finite and naturally exhausted as the trading period progresses (see, for instance, the classical experiments described by Smith (1962) or Plott (1982)), in the CRSD environment there is no natural end to trading.

Brewer, et al. (2002) describe the particulars of the CRSD environment as follows:

> "... if buyer #3 used a private market offer (a redemption value) from the experimenter, this same offer would immediately be made to the *next* buyer (e.g., buyer #4). Similarly, offers to sell (costs) were recycled to the next seller. Subjects had no knowledge at all about this refreshing. Subjects knew only that new orders could appear in their private markets at any time.
>
> Refreshing the private offers in this way keeps the instantaneous supply and demand curves constant at every moment in time. If an offer is used or expires, it does not vanish from the pool of supply and demand. Instead, it is recycled to someone else. Thus, the opportunities of gains from trade are never exhausted. The market demand and supply functions as represented by redemption values and costs are always constant – independent of the patterns of trade."

Figure 2 shows the data set of public market trading prices produced from 2½ hours of trading in the environment of Figure 1. The data shown here has been 'sanitized' by removing possible outliers or errors – trades with large price movements – and will serve as the primary data source for this paper.

There are three primary benefits of the CRSD environment over other sources of data: (i) CRSD can produce long time series (there are 793 trades in the sample we will use vs. ~20 in the typical double auction period) useful when examining time series properties as the accuracy of some of the related estimators scales only as $1/\sqrt{N}$; (ii) because of the nature of the refreshing, the instantaneous supply and demand is held stationary; one does not have to consider the possibility of an equilibrium price that is changing as traders exit the market; (iii) stationary instantaneous supply and demand can be useful in separating models of market behavior and convergence. Certain price convergence processes – such as Marshallian path processes and noisy analogs like the Gode and Sunder (1993) ZI Robots – that operate

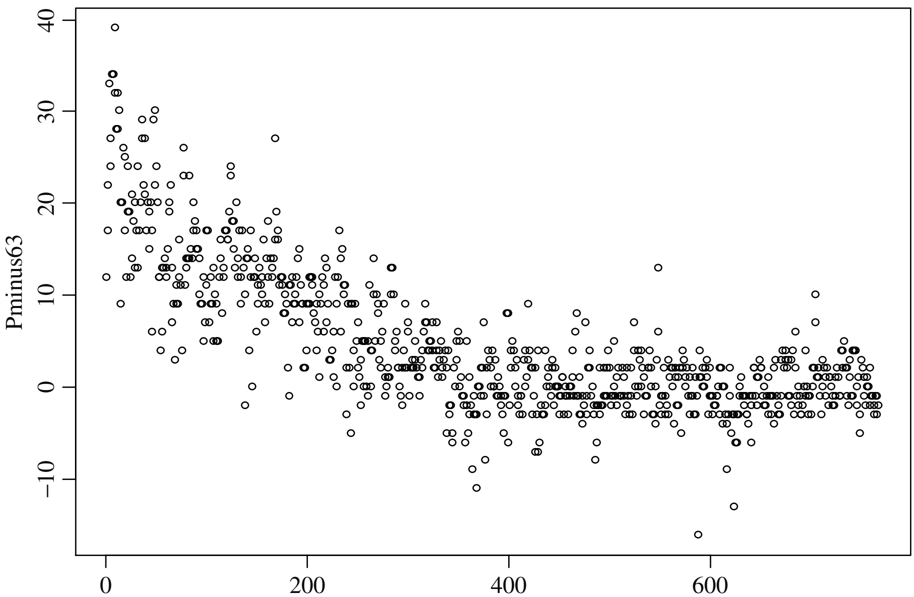

Figure 2. Price Time Series P1 from Brewer, Huang, Nelson, and Plott (sanitized).

in the ordinary double auction can not operate in the CRSD environment and therefore can not be an explanation for why human-populated CRSD markets are observed to converge to an equilibrium price.

3. STANDARD MODELS

Currently, fundamental models of market processes differ somewhat in both form and function between the fields of microeconomics and finance. The purpose of this section is to illustrate these basic models – much of which may be quite familiar to some readers. Section 4 will then consider how these models overlap in ways that might be compatible or incompatible.

3.1. The Microeconomics Approach – Law of Supply and Demand, Allocation Efficiency, and Dynamic Adjustment

The Law of Supply and Demand is a static theory of market equilibrium, and provides that the equilibrium of a competitive market occurs at the price and quantity given by the intersection of the demand and supply curves. For example, in Figure 1, the intersection of the demand and supply curves gives $Q = 6$ and $55 \le P \le 60$.

In an ordinary market experiment without the continuous refreshing described in section 2, the equilibrium $Q = 6$ and $55 \le P \le 60$ would be the predicted outcome of microeconomic competitive theory. With continuous refreshing of the supply/demand curves, the correct equilibrium concept is less clear. First, there is no prediction for Q because refreshing allows trade to continue. Brewer, et al. (2002) maintained a hypothesis that competitive theory could possible still predict prices in these environments: the prediction for P from the Figure 1 of $55 \leq P \leq 60$ is the "instantaneous competitive equilibrium" of Brewer, et al. (2002), and there may also be a velocity-based equilibrium based on observed supply and demand curves that are calculated ex-post.

Allocation efficiency measures the ratio of the gains from trade achieved in a market versus the maximum possible gains from trade. Ordinarily, laboratory markets in the absence of externalities will equilibrate to prices near those predicted by the Law of Supply and Demand, supporting an allocation having nearly 100% allocation efficiency. In continuously refreshed markets, allocation efficiency is difficult to define because the proper definition of maximum possible gains from trade in the presence of refreshing is not obvious.

The Law of Supply and Demand is not a dynamic theory of price adjustment. Two early models of price adjustments are due to Marshall and Walras. Either could be expressed in the form of a differential equation, though there is no exact differential equation known to be accepted as an exact realization of either theory. The primary difference between the two adjustment models is whether adjustment occurs along the quantity axis or the price axis.

3.2. Marshallian Adjustment

The Marshallian adjustment process can be written as:

$$\mathrm{d}Q/\mathrm{d}t = F(P_D(Q) - P_S(Q))$$

where

$P_D(Q)$ = Demand Price (or Marginal Value) at Q

$P_S(Q)$ = Supply Price (or Marginal Cost) at Q

$F()$ is a sufficiently well-behaved unknown monotone function

3.3. Walrasian Adjustment

The Walrasian Adjustment Process can be written as:

$$\mathrm{d}P/\mathrm{d}t = G(Q_D(P) - Q_S(P))$$

where

$Q_D(P) - Q_S(P)$ gives the quantity of the excess demand at price P

and $G()$ is a sufficiently well behaved unknown monotone function.

The Marshallian adjustment process was originally associated with adjustment of markets that are repeated over a series of days, months, or years. One general argument is that if $P_D > P_S$ it will be easy for sellers to sell their goods in excess of their marginal cost, and production will expand. However, if $P_D < P_S$ trade will be difficult since buyers are willing to pay less than sellers require to meet production costs. Because some sellers will be producing at least part of their production at marginal costs that are higher than what buyers are willing to pay (P_D), sellers must necessarily take a loss on this excess production. Failure to sell at a price greater than marginal cost would rationally lead to a contraction of production over time as sellers learn to correct overproduction.

Walrasian adjustment can be thought of as either a virtual or real tatonnement process that occurs before trade to set a price, or it can be thought of as occurring within trade through shortages and surpluses. In this research, we are mainly concerned with the latter approach. The basic idea is that if prices are above equilibrium, there is excess supply, and prices will fall over time, and if prices are below equilibrium there is excess demand, and prices will rise over time. It is worth noting that the Walrasian adjustment process, as a first-order differential equation, implies an exponential approach to equilibrium. A first-order differential equation for price over time does not permit more advanced behavior seen in some physical (non-economic) systems: for example, the oscillation of a spring (with or without damping) is the result of a 2^{nd} order differential equation of motion over time.

More recently, Easley and Ledyard (1993) provide a model of double auction price convergence that has both Marshallian and Walrasian aspects. However, this model applies to the standard double auction with finite periods, not the CRSD double auction.

Attempts at comparing the Walrasian and Marshallian adjustment processes in standard double auctions have been made by Plott and George (1992) and Jamison and Plott (1997). These studies involved the creation of externalities alternatively generating upward sloping demand or downward sloping supply (called "perverse-shaped" curves because normally demand is downward sloping and supply is upward sloping) to create particular regions of Walrasian instability/Marshallian stability or Marshallian instability/Walrasian stability. Plott (2001; Introduction p. xxv) summarizes these results as favoring a Marshallian theory when externalities cause perverse-shaped supply and demand curves but favoring Walrasian theory when income effects cause perverse-shaped curves.

One can see that there is no well-accepted choice between Marshallian and Walrasian dynamics. It is believed that the use of the Continuously Refreshed Supply and Demand in the research reported here will select against Marshallian dynamics because there will be no shortage of trading opportunities along the Q axis to force an outcome. This consequence of CRSD experiment design will be revisited again in the next section.

3.4. The Financial Economics Approach – Informational Efficiency

Market prices are said to be *informationally efficient* if prices summarize existing information to the extent that there is zero expected gain from buying or selling based on existing information. Existing information includes all current and prior prices $\{P_t, P_{t-1}, P_{t-2}, \ldots\}$ as well as any other commonly known information about the market.

More formally, given information I_t at time t, prices are a *Martingale process* whereby the expectation $E_t[P_{t+k}|I_t] = P_t$ for all $k > 0$.

3.5. Normal Random Walk

For the purpose of this paper, a normal random walk is an integrated time series P_t whose first differences $\Delta P_t = P_{t+1} - P_t$ are independently and identically distributed normal variables with $E(\Delta P_t) = 0$ and $\mathrm{Var}(\Delta P_t) = \sigma^2$. Modeling prices in a market as a random walk necessarily satisfies the informational efficiency requirement: if the mean of the difference process is zero, then $E[P_{t+k}|P_t] = P_t + E(\Delta P_t) + E(\Delta P_{t+1}) + \ldots = P_t + 0 + 0 + \ldots = P_t$. Note that the normal random walk has linearly increasing prediction variance $\mathrm{Var}[P_{t+k}|P_t] = k\sigma^2$ as the prediction horizon k is increased.

3.6. Heteroskedastic Martingales

A heteroskedastic Martingale is a time series that satisfies the informational efficiency hypothesis but is not a normal random walk due to changes over time in the variance parameter of price differences σ^2. The variance could be time dependent or price dependent. Well known examples of this class of processes would include the ARCH and GARCH time-series models, which add a separate equation for variance that induces heteroskedasticity.

4. COMPARING AND CONTRASTING THE STANDARD MODELS

The financial and microeconomic theories appear to overlap only in the case of a perfect market that instantaneously finds the competitive equilibrium price. A constant price is trivially a Martingale and if this constant price is at the theoretical equilibrium then both kinds of theories can be satisfied. However, noisy prices are an empirical regularity common to both the lab and the field. The main tension between the two approaches of Section 3 is that the existence of a price adjustment process in Microeconomics converging towards the static prediction of the Law of Supply and Demand is incompatible with the notion that markets prices exhibit informational efficiency detectable through autocorrelation properties of price differences.

4.1. Random Walk destroys convergence

If prices were a random walk, the market would have informational efficiency but then prices would not converge towards any fixed level. Price increments are always independent and identically distributed and therefore do not tend to move price towards the microeconomic competitive equilibrium given by the Law of Supply and Demand.

4.2. Convergence of prices towards competitive equilibrium implies non-Martingale behavior

Convergence of prices towards a competitive equilibrium price p^* would seem to suggest that $E_t[P_{t+k}|I_t] < P_t$ when $P_t > p^*$ and $E_t[P_{t+k}|I_t] > P_t$ when $P_t < p^*$. In contrast, a Martingale Process always has expectation $E_t[P_{t+k}|I_t] = P_t$ for all $k > 0$.

Voluntary trade within a set of supply and demand curves necessarily generates a price ceiling and a floor outside of which trade will never occur. This creates problems for random walk and Martingale models.

4.3. Voluntary Trade and The Support of Possible Prices

Nothing in a random walk theory prevents prices from wandering outside of the support of voluntary trade. For example, in Figure 1 the lowest seller's marginal cost is 30 and the highest buyer's marginal value is 140. Voluntary trades can only occur at prices greater than or equal to 30, and less than or equal to 140.

4.4. A Censored Normal Random Walk is no longer a Martingale process

Censoring the random walk above and below ceiling and floor values (P^H, P_L) would tend to violate the Martingale requirement that $E[P_{t+k}|P_t] = P_t$. To see this, consider a price ceiling P^H, then at $P_t = P^H$ we would necessarily have $E[P_{t+1}|P_t] < P_t$. Unless $P_{t+1} = P_t = P^H$ with certainty (which is never true for a censored iid normal random walk but could be true for a heteroskedastic censored random walk only for the unusual case that the variance falls to zero at the ceiling) the mean of the next price P_{t+1} must be less than the ceiling because the probability support does not include any prices above the ceiling. The argument for violation at a floor is similar.

4.5. Bounded Martingales seem to require various non-economic properties

A Bounded Martingale is a Martingale price process bounded between two limits $[P^L, P^H]$. From the previous paragraph we know that the first non-economic property that a Bounded Martingale must have is that the price bounds are sticky. If at some time t the price $P_t = P^H$ or $P_t = P^L$, it remains at P^H or P^L forever. If one considers exponentially decreasing, ever-tightening bounds on the variance of the Martingale process over time, one may obtain price convergence to an interior point, but there is no reason to believe that this interior point should always coincide with the economic notion of competitive equilibrium nor does classical economics provide a definitive source or model associated with this decrease in variance. Conditional heteroskedasticity based upon the distance of the price P_t from equilibrium might be helpful, but once again there is no obvious economic source of this effect and one must still have a variance of 0 at P^L and P^H with the possibility of prices becoming stuck at or near these locations. In contrast, economic theory would seem to say that the forces pushing prices towards the equilibrium would be strongest at boundary prices P^L and P^H because it is at these prices that excess demand or excess supply will be greatest.

The next note about conflicts among the models has more to do with the specific choice of a CRSD environment for generating the experimental data.

4.6. CRSD environment selects against Marshallian Price Adjustment Processes

Brewer, et al. (2002) considered an alternative interpretation of the Marshallian adjustment process acting within a single trading period: the *Marshallian Path*. The idea is simply that the sequence of trades in a market will be from left to right along the supply and demand curves at any series of prices P_n where $P_S(n) \le P_n \le P_D(n)$. For example, for the market of Figure 1, the Marshallian Path theory would imply the following sequence of trade: (buyer with value 140/seller with cost 30), (buyer with value 125/seller with cost 35), (buyer with value 110/seller with cost 40), (buyer with value 95/seller with cost 45), (buyer with value 80/seller with cost 50), (buyer with value 65/seller with cost 55). The equilibrium quantity of trades would be $Q^* = 6$. No further trades will be possible since $P_D < P_S$ at $Q = 7$.

Gode and Sunder (1993) advanced the idea that fully human rationality suggested in the adjustment processes above was not necessary because markets populated by so-called "Zero Intelligence" robots, which patiently bid/ask randomly within their budget constraint, converged to market equilibrium prices. ZI robots effectively follow a noisy Marshallian path, because at any time the robots with the greatest probability of trading are the high value buyer and the low cost seller. By removing the high-value and low-cost traders early, prices are stochastically forced towards the competitive equilibrium at the supply-demand intersection. Cason and Friedman (1996) provide additional evidence that suggests markets populated by humans follow such a noisy Marshallian path.

The continuously-refreshed environment of Brewer, et al. (2002) removes the Marshallian path as a possible mechanism for adjustment because the high-value and low-cost units are recycled back into the market. Prices are still seen to converge. This might be seen as lending support towards a Walrasian adjustment model at least for the CRSD class of environments.[2]

5. A HYBRID MODEL – ROBOT SIMULATIONS

Figure 3 shows market prices generated by three groups of specially designed trading robots. These prices are seen to converge towards a kind of equilibrium, similar to the convergence of the humans. The robots, which we will call constrained random walkers, use a pricing algorithm based partially upon a random-walk. The purpose of

Figure 3. Prices from Constrained Random Walkers attain equilibrium over time.

this section is to explain the algorithm of these robots, compare this algorithm to the Zero Intelligence algorithm of Gode and Sunder (1993), and compare and contrast the behavior of markets populated by the robots. The potential significance of these robot simulations for a combined microeconomic/financial theory of markets is explored.

5.1. Constrained Random Walkers

Constrained Random Walkers obey the following algorithm: at each moment in time one robot representing a particular buyer or seller is selected to act. This robot will then (1) fetch the previous transaction price p_{t-1}. This transaction price is the price of the last completed trade, not the advertised price of a previous bid or ask. (2) add an independent, identically distributed deviate $\varepsilon \sim N(0, \sigma^2)$ to obtain the potential price $p^* = p_{t-1} + \varepsilon$, and (3) submit a bid or ask at price p^* if and only if p^* is within the robot's budget constraint – that is, $p^* >$ cost for a seller, or $p^* <$ value for a buyer. Potential prices that fail step (3) are discarded.

5.2. Gode and Sunder (1993) ZI Robots

The ZI Robots obey the following algorithm: at each moment in time one robot representing a particular buyer or seller is selected to act. This robot will then randomly bid over the budget constraint without regard to previous prices. A Buyer robot will bid a price b^* from a uniform random distribution over $0 \leq b^* \leq v$, where v is the redemption value. A Seller robot will ask a price a^* from a uniform random distribution over $c \leq a^* \leq H$, where c is the cost of the unit to the seller and H is an arbitrary ceiling price chosen to be higher than the highest buyers' value.

5.3. Double Auction Trading Rules

As bids and asks arrive, they are interpreted under the two rules of the double auction. The first rule is an improvement rule that discards bids and asks if they are not better than any previous standing bid or ask. The second rule is a trading rule that specifies that a trade occurs when a new bid is greater than the ask price, or a new ask is less than the bid price. When a trade occurs, the earlier bid or ask of the pair determines the trading price.

5.4. Effect of Individual Budget Constraints

The effect of individual budget constraints manifesting the supply and demand curves must be significant for any organized trend of prices towards an equilibrium price predicted by the Law of Supply and Demand. Gode and Sunder (1993) demonstrate that without the individual budget constraint and the double auction improvement rule requiring bids to be ascending and asks to be descending, the ZI robots do not converge to an equilibrium price but instead generate independent, identically distributed prices over the interval $[0, H]$. Brewer, et al. (2002) demonstrated the additional requirement of scarcity or finiteness of supply and demand for the ZI robots to reach equilibrium prices. In the CRSD environment, ZI robots fail to reach an equilibrium price, instead generating an iid sequence of prices. However, the exact shape of the iid distribution is affected by the particulars of the supply and demand curves.

Without individual budget constraints, the Constrained Random Walkers would generate prices that are a Martingale process. The individual budget constraints are imposed at step (3). Without step (3), each proposed bid or ask price is simply a normal based around the previous price. But with step (3) added, prices appear to converge. It is clear from Figure 3 that the rate of convergence depends on the deviation parameter σ^2 of the Normal distribution generating successive bids and asks. Over a range of small σ^2, higher σ^2 appears to allow convergence to proceed at a faster pace.

5.5. Effect of Double Auction Trading Rules

The effect of the double auction trading rules is to impose a type of order on the competitive process that converts streams of bids and asks into transaction prices. The importance of these rules, and of changes to them, is borne out by the rich literature of double auction processes. Chamberlin's (1948) experiments showing the apparent non-convergence of market prices did not impose the formalities of double auction trading, but instead had subjects circulate the room to find partners. In contrast, Smith (1962) showed that when the rules of the double auction were applied to trading, prices converged after a series of repetitions to match the predictions of the Law of Supply and Demand.

5.6. Price Convergence in CRSD markets populated by Constrained Random Walk Robots

Brewer, et al. (2003) showed that with a CRSD environment, the ordering effect of the double auction market lacks sufficient strength to tame the aggregate pricing behavior of the ZI robots. However, because prices of markets populated by humans converge in the CRSD double auction environment, it was hypothesized that some additional element of human rationality, absent in the ZI robots, was responsible. With the demonstrated convergence of market prices in double auctions populated by the Constrained Random Walkers, the element of behavior required may have been identified: basing of bid and asks upon the previous price, while still censoring bids and asks against the budget constraint, causes the market prices to converge.

Notice what happens as the robots compete in Figure 3. Prices drift towards equilibrium at a rate that rises with increasing innovation σ^2. After noting the pattern, σ^2 was varied in an attempt to generate time series comparable to the human traders. But why should prices converge at all? The key is to recognize the combined effect of budget constraints, double auction rules, and anchoring bids and asks to the previous transaction price.

If the previous price is low compared to competitive equilibrium, then the budget constraints imply a larger pool of buyers submitting bids than sellers submitting asks. The double auction rules require bids to be ascending and asks to be descending. Suppose prices are so low that it is likely that 2 buyers will submit bids before the next seller will submit an ask. Then the double auction rules will filter out the highest of these two bids, which has a 75% probability of being higher than the previous transaction price. While the bid price will likely move up, it is unlikely it moves up by much more than σ because of the anchoring effect of the bid generation process where $b \sim p_{t-1} + N(0, \sigma^2)$. When the seller robot generates an ask, with about 50% probability the ask price will be below the previous transaction price and a trade will occur at the earlier, and higher, bid price. Therefore the trade price will tend to move slowly towards the equilibrium, with the strength of the drift decreasing as prices move towards the equilibrium. When the prices are too high, there are more potential sellers than potential buyers, and a similar process occurs to lower the price.

This type of slow convergence suggests an AR(1) process might reasonably fit prices converging towards competitive equilibrium, compatible with the notion of Walrasian adjustment processes:

$$(P_{t+1} - P_{eq}) = a_1(P_t - P_{eq}) + \varepsilon_t;\ |\,a_1\,| < 1,\ \varepsilon_t \text{ iid } N(0, \sigma^{*2})$$

The market prices of the Constrained Random Walkers fit an AR(1) process fairly well. It is possible that there could be some price-based heteroskedasticity that does not fit the standard AR(1) model, or the residuals may be non-normal. These effects were not tested formally. When we look at the data of the human populated markets, there is also an additional effect that does not fit a AR(1) process: correction of outlier prices. The analysis of markets populated by humans will be the focus of the next section.

6. ARMA BEHAVIOR OF MARKETS POPULATED BY HUMANS

The purpose of this section is to examine ARMA models of the CRSD double auction market populated by humans. The impetus for using ARMA models is based in part upon the hypothesis that markets populated by Constrained Random Walker robots of section 5, which demonstrate convergence towards competitive equilibrium, appear to fit an AR model in prices.

However, with the humans, the visual evidence suggests a handling of outliers inconsistent with a simple AR(1) model. In an AR(1) model, an outlier in price would generate a new slow drift towards the equilibrium price. But in this data, the observed effect is that the price corrects to a price near the previous prices. This is a property of a moving average or MA model where the error terms follow a linear process and allow for such self-correction. An ARMA(1, 1) model incorporates both effects.

$$(P_{t+1} - P_{eq}) = a_1(P_t - P_{eq}) + \varepsilon_t;$$
$$\varepsilon_t \sim b_1\ \varepsilon_{t-1} + \text{iid } N(0, \sigma^{*2})$$

In this model, the a_1 term is typically denoted the AR(1) or autoregressive term and the b_1 term is typically denoted the MA(1) or moving-average term. Slow convergence towards equilibrium is described by a near unity $a_1 \sim 1 - \phi$, with the speed of convergence increasing with ϕ. The b_1 term indicates "memory" in the shocks. A positive b_1 may indicate a run-on effect in large shocks being followed by a run of smaller and smaller shocks. A negative b_1 may indicate that shocks tend to partially self-correcting in successive trades. From a visual inspection of the human trading data, we expect b_1 to be negative in human populated markets.

The analysis of the data yields six results. Result 1 states that neither a fixed random process nor a random-walk unit-root process adequately describes the human market data. Result 2 identifies the drift in the pricing process and identifies a large source of variance from outliers, or large movements in price that are almost immediately corrected[3]. Based on this, we removed large movements in price to "sanitize" the time series. The goal is to separate the effects of these self-correcting price movements from other features of the time series. Result 3 finds a curious relationship between price variance and price in the sanitized time series. Results 4–6 characterize features of ARMA models fitted to the time series.

Result 1: Neither an iid fixed random process nor a unit-root process – such as a random walk – adequately describes the price data.

Support: Visually, it is unlikely that the data could be independent and identical draws from a fixed random distribution because the mean and variance of the process are changing. Visually, a unit-root process is unlikely because shocks to a unit-root process are persistent. This means, for instance, that large changes in the price should not be followed by reversals. Two formal tests were performed to examine the possibility of a unit-root. Both the Dickey-Fuller test and the Philips-Perron test for a unit-root yield p-values of less than 0.01 for this data, indicating rejection of the null hypothesis of a unit-root at 99% certainty.

Trades	1–100	101–200	201–300	301–400	401–500	501–600	601–700	701-end
Raw Data								
Mean	81.22	75.86	69.89	64.83	63.44	63.39	61.91	63.40
Var	87.11	26.53	21.49	18.06	18.03	13.78	9.62	6.61
$\lvert \Delta P \rvert > 15$ removed (26 trades)								
Mean	80.69	74.82	68.88	64.14	62.85	63.18	62.05	63.47
Var	57.87	24.96	22.51	15.03	9.85	9.93	8.03	6.58

Result 2: The price data shows a slow drift in mean from $T = 1$ to $T = 300-400$. The variance also changes over time, and is generally decreasing. A large portion of the variance in trades 1–100, 401–500 and 501–600 can be attributed to price changes where $\lvert P_{t+1} - P_t \rvert > 15$.

Result 3: The variance of the price process appears to be exponentially decreasing, once certain outliers are removed.
Support: A naïve OLS regression of log(var) on the group midpoints T_{mid} = (50, 150, . . .) of groups of 100 observations yields

$$\log(\text{var}_{100}(p_t)) \sim \underset{(\pm 0.159)}{3.8379} \quad \underset{(\pm 0.00034)}{-0.00283\, T_{\text{mid}}}; \qquad \text{(standard errors)}$$

The adjusted R^2 of this model is 0.9041 indicates a fairly close match as can also be seen visually in Figure 4.

The price data exhibits some features of an ARMA(1, 1) model, provided one is willing to ignore the heteroskedasticity and the possibility of higher order terms.

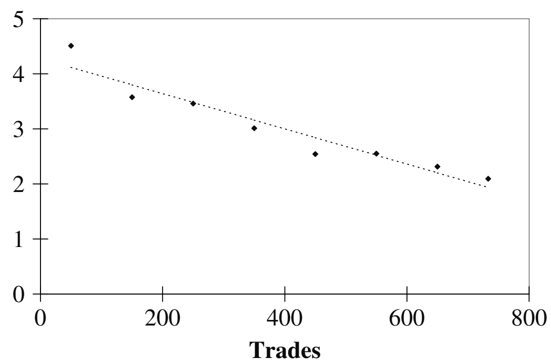

Figure 4. Log(Var[P]) for Groups of 100 trades with Log-linear Fit.

ARMA (1, 1) model		AR(1)	MA(1)	σ^2	Log(Likelihood)
Raw $P[t]$ no filtering		0.9952	−0.8043	23.73	–2382
	(se)	0.0041	0.036		
$\lvert\Delta P\rvert > 15$ removed		0.9877	−0.6429	15.95	–2151.6
	(se)	0.0067	0.0419		
	AR only	0.8726	–	19.80	–2234.0
	MA only	–	0.6567	42.44	–2526.0

Result 4: The ARMA(1,1) fits reveal (i) an AR coefficient compatible with a very slow Walrasian dynamic together with (ii) a stronger MA coefficient compatible with short-term corrections of remaining outliers against the slowly moving mean. Sanitizing the data enables better detection of the slow Walrasian dynamic.

Support: The strength of the convergence process depends on $(1 - a_1)$. As the a_1 coefficient is almost 1.0, the convergence process is very slow and furthermore does not have good statistical significance given the standard error of the a_1 coefficient. The MA(1) coefficient b_1 is negative and is picking up the bounce or correction of large movements in price. Removing the large price changes from the time series improves the log likelihood by over 200 and shows a slightly stronger convergence dynamic now safely above the noise. The AR(1) and MA(1) process estimated separately show that both terms are significant. A log-likelihood χ^2 test would reject removing either term at well above the 0.999 level.

Result 5: A structural change in the ARMA process may occur roughly corresponding to the attainment of equilibrium.

Support: Figure 5 shows a standard log-likehood test for detecting the breakpoint for a single structural change in a time series model. Figure 5 suggests, based on log-likelihood, a structural break around $T \sim 290$. When we look at the time series of prices this does correspond to a rough visual assessment of where equilibrium appears to have been attained ($T \sim 300–400$).

		Coefficients			
ARMA 1, 1 models $\lvert\Delta P\rvert > 15$ removed		AR(1)	MA(1)	$\sigma 2$	Log(Likelihood)
$T \leq 290$		0.9882	−0.5918	26.02	−885.25
	s.e.	0.0088	0.0631		
$T > 290$		0.7368	−0.4615	9.05	−1204.7
		0.0846	0.1138		
				Combined	−2089.95

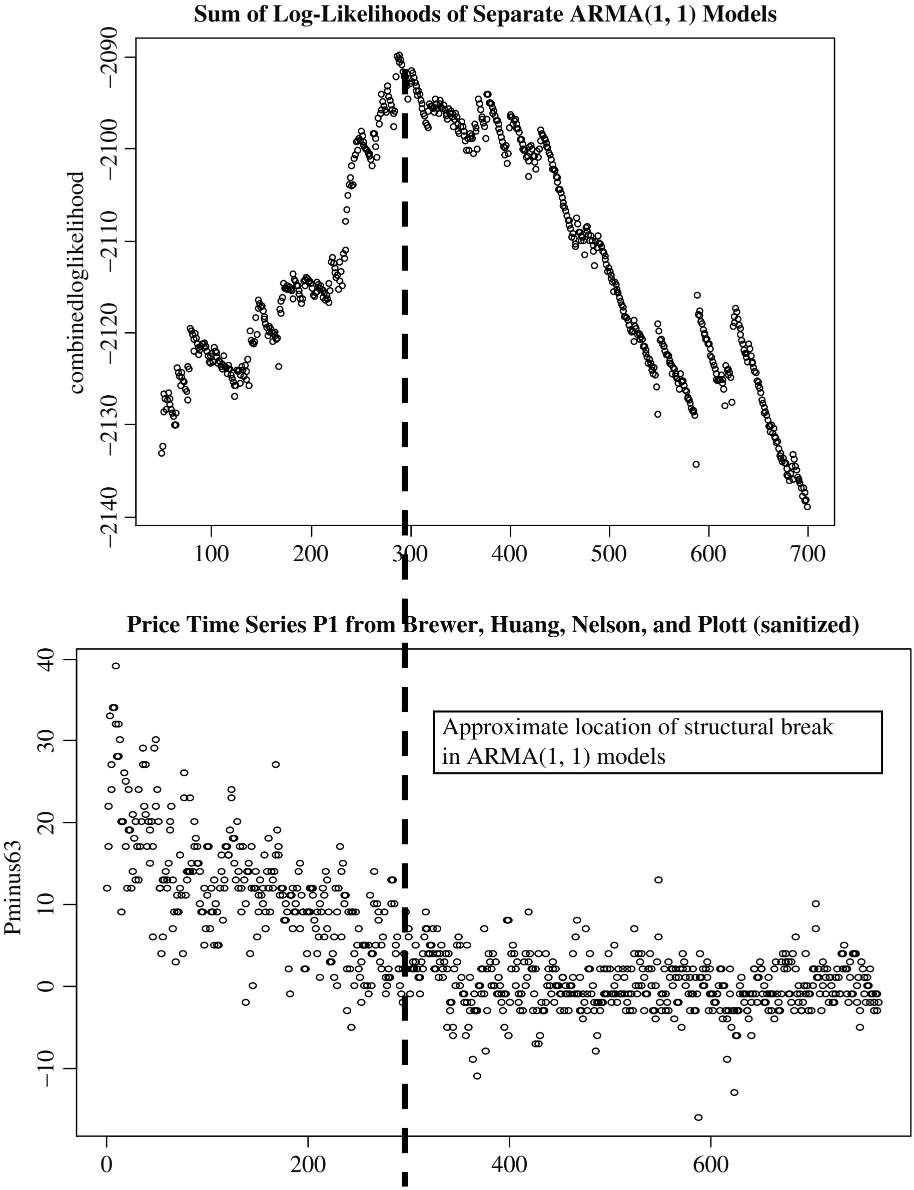

Sum of Log-Likelihoods of Separate ARMA(1, 1) Models
combinedloglikelihood
−2090
−2100
−2110
−2120
−2130
−2140
100
200
300
400
500
600
700
Price Time Series P1 from Brewer, Huang, Nelson, and Plott (sanitized)
Pminus63
40
30
20
10
0
−10
0
200
400
600
Approximate location of structural break
in ARMA(1, 1) models

Figure 5.

With the caveat that smaller data sets yield less accurate fits, it is possible to break up the time series into groups of 100 trades.

		Coefficients		
ARMA 1, 1 models	AR(1)	MA(1)	$\sigma 2$	Log(Likelihood)
Data				
all data, no filtering				
trades 1–100	0.9954	−0.9085	78.98	–
101–200	0.996	−0.7463	24.65	–
201–300	0.9789	−0.6892	20.11	–
301–400	0.857	−0.4291	12.35	–
401–500	0.2489	−0.1507	22.6	–
501–600	0.1925	0.0329	8.53	–
601–700	0.8711	−0.6732	9.23	–
701–792	0.4172	−0.099	5.99	–
$\lvert\Delta P\rvert > 15$ removed				
trades 1–100	0.983	−0.4676	35.44	−321.97
101–200	0.9907	−0.676	20.33	−293.68
201–300	0.962	−0.5767	19.84	−291.97
301–400	0.9058	−0.6693	12.23	−267.30
401–500	0.674	−0.4279	8.49	−248.91
501–600	0.0083	0.1263	9.68	−255.42
601–700	0.7797	−0.422	6.68	−237.00
701–767	0.5295	−0.179	5.68	−153.34
			Combined	−2069.59

The pattern is given as Result 6.

Result 6: The price convergence process is not purely stationary but there is (i) a slow trend towards lower variance, and (ii) a shift in behavior around attainment of price equilibrium toward stronger price-related effects and away from outlier effects.

Support: The variance effect is clearly shown in the table above, and is also expected given Result 3. There are large changes in the coefficients beginning with $T = 301–400$. Figure 6 (constructed from the unfiltered data) suggests that the AR(1) coefficient drops, indicating a stronger strength of the price convergence process. At the same time, the MA(1) coefficient drops in absolute value, indicating that less of the price movements seem to be corrections of previous shocks.

Note also that the coefficients begin to wander after $T > 500$. This wandering may involve effects of discreteness in prices, given that prices are constrained to unit values and the observed price variance is very low.

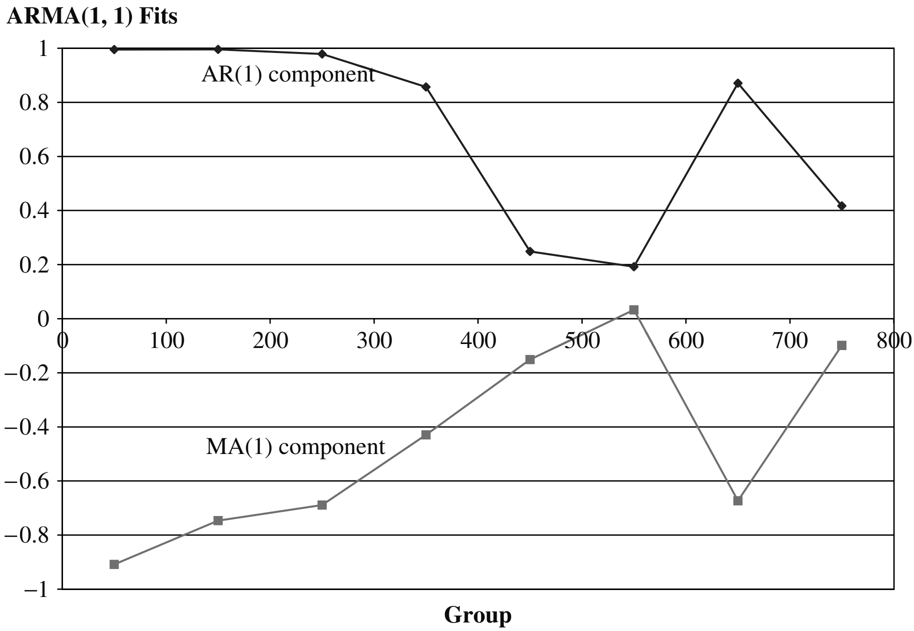

Figure 6. ARMA(1, 1) fits to raw subsamples.

An alternative, but perhaps important, interpretation for Result 6 is that as the market equilibrates, there is initially an advantage for subjects who pay attention to the trend in market prices and whether a price offer is an outlier and this generates the near unity values for the time-series coefficients that we see. But as the market equilibrates, there is less and less variance in price and therefore less and less money that could be earned by careful timing of, or attention to, market activities. The changes and variability in the time-series properties that occur upon, or just after equilibrium, could be simply due to the fact that there is almost nothing to be gained by trading in the previous manner.

7. CONCLUSIONS

This paper began with the question of whether it might be possible to integrate or reconcile ideas of market dynamics found in microeconomics with those found in the random walk or Martingale theory of finance. The use of long time series generated by a CRSD laboratory market provided a practical framework for an initial study of these questions.

Two conclusions can be found in the present research. The first conclusion is that something like a random walk process can be useful in modeling the slow convergence component of prices found in CRSD markets. When a random walk in bids and asks is censored against individual budget constraints the resulting market prices appear to slowly converge towards the predictions of supply and demand. The innovative step in this model is that the random walk is not in transaction prices, but instead is a component involved in the process generating bids and asks.

The second conclusion is that the price dynamics of human-populated markets contain a number of different kinds of effects that seem to be operating simultaneously. Smoothing shows a AR(1) process similar to that seen in the constrained random walk robots. However, prices in the human-populated markets also show a complex outlier generation and correction process. A large move in prices at one trade is often corrected back towards the average with the next trade. This type of "memory" of the process is not captured by an AR(1) statistical process or a constrained random walk of bids/asks. Removing many of the large outliers and adding an MA(1) component to absorb the remaining outlier/correction process yields an ARMA(1, 1) model that varies as the market converges towards equilibrium.

A structural break in the ARMA parameters seems to occur as equilibrium is reached. The nature of this structural break is left for further research. It may suggest the use of models with multiple regimes for price discovery and equilibrium behavior rather than a simple stationary model.

Hope for a combined theory of microeconomic and financial adjustment suggested in this work possibly relies on classifying markets along what is for now a somewhat speculative grouping: (i) Markets with finite ending times and finite trade can be roughly modeled as a noisy Marshallian process and it is possible that scarcity and the likelihood of large surplus traders trading early are all that is necessary for the appearance of prices converging to the competitive equilibrium. Prices following a noisy Marshallian process will not be Martingales but instead will exhibit negative autocorrelation of price changes. (ii) Markets with no fixed ending time and continuously refreshed supply and demand, such as the CRSD market presented here, exhibit price convergence when populated by humans that can not be explained as a Marshallian process, but only as either a Walrasian class of adjustment process or some other type of process yet to be defined. Within the Walrasian class of adjustment processes is the possibility that a random-walk approach to submitting bids and asks can explain market price convergence when the random walk generating bids and asks is censored, at the individual level, against an individual seller's supply costs or buyer's redemption values making up the supply and demand. This convergence also relies upon the details of trading, e.g., the improvement rules of the double auction.

It is this second grouping of markets involves a hybrid model of CRSD market price convergence incorporating both ideas from finance (the random walk) and microeconomics.

In addition, the hybrid model can yield back a purely financial model, when the microeconomic constraints of supply and demand are removed. Suppose a financial market is modeled as being like the CRSD environment, but without the individual constraints on traders' costs and values typically associated with supply and demand. If individual behavior based on (possibly faulty) future expectations turns out to be quite different from behavior based on a known arbitrage opportunity, this may be a sensible rough model. In the absence of individual budget constraints associated with supply and demand curves, a random walk generating bids and asks would be uncensored, depending only on the previous transaction price. Thus, when the rules of double auction trading are applied, the random walk process in bids and asks would generate a random walk in prices, which is the expected result in a informational efficiency model of a price adjustment.

As warned in the introduction, this approach may be seen by some as overly broad or raising more questions than it answers. However, as pointed out in section 4, there are a number of apparent conflicts between microeconomic price processes as confirmed by laboratory experiments and the standard assumptions of finance. Rather than simply assume that the latter do not apply to the former, it is hoped that the search for a combined theory of price adjustment may – with continued contribution by theorists and empiricists in both fields – yield further insights into the behavior of markets not discernable with the tools of one field alone.

ACKNOWLEDGMENT

This research was supported in part by Hong Kong Research Grants Council Competitive Earmarked Research Grant HKUST6215/00H. I also wish to thank Tim Cason, Leonard Cheng, S. H. Chew, Jim Cox, John Dickhaut, Dan Friedman, Steve Gjerstad, John Ledyard, Charles Noussair, Charles R. Plott, Jason Shachat, and Shyam Sunder for their comments on previous installments of this research: none are responsible for my errors.

NOTES

[1] Fama [1965] finds positive autocorrelations of returns on 75% of the Dow 30 stocks. Solnik [1973], Lawrence [1986], and Butler and Malaikah [1992] find many examples of negative autocorrelations of daily returns in the stock markets. Hawawini and Keim [1995] provide a survey of this literature in their introduction. Dacorogna, Gençay, Müller, Olsen and Pictet [2001; pp. 123–4] briefly discuss negative autocorrelations found in the price formation process of foreign exchange markets that are strong for very short time horizons but vanish over the course of 30 minutes or so.

[2] The argument is that in a CRSD environment, the Marshallian dynamic is precluded. Nothing precludes Marshallian or multiple adjustment models from acting in the ordinary classical environment or other domains not studied here.

[3] The interface for trading was a graphical screen rather than numeric input of bids and asks, and so it is also possible that these outliers are caused by subjects recklessly sending in asks that are too low (or bids that are too high) in an attempt to quickly accept a desirable bid (ask). Priority for input from multiple subjects that would cause acceptance of a bid/ask was based on time only not on price. A tendency for some subjects to send in outlier prices can also be seen as a human/computer interface design flaw because subjects could have been warned and given a pop-up box to confirm questionable behavior. The design used did not involve any such handholding, settling for simplicity and minimal interaction or guidance of the subject.

REFERENCES

Bossaerts, Peter L. (2002). *The Paradox of Asset Pricing* Princeton: Princeton University Press.

Brewer, Paul J., Maria Huang, Brad Nelson and Charles R. Plott (2002). "On the Behavioral Foundations of the Law of Supply and Demand: Human Convergence and Robot Randomness." *Experimental Economics* 5: 179–208.

Butler, K. C., and Malaikah, S. J. (1992). "Efficiency and inefficiency in thinly traded stock markets: Kuwait and Saudi Arabia." *Journal of Banking Finance* 16, 197–210.

Cason, Timothy N. and Daniel Friedman (1993). "An Empirical Analysis of Price Formation in Double Auction Markets," in D. Friedman and J. Rust (eds.) *The Double Auction Market: Institutions, Theories, and Evidence* pp. 63–98.

Cason, Timothy N. and Daniel Friedman (1996). "Price Formation in Double Auction Markets." *Journal of Economic Dynamics and Control* 20, 1307–1337.

Chamberlin, E. H. (1948). "An Experimental Imperfect Market." *Journal of Political Economy* 56(2), 95–108.

Dacorogna, Michel M., Gençay, R., Müller, U., Olsen, R. B., and Pictet, O. V. (2001). *An Introduction to High Frequency Finance* San Diego: Academic Press.

Duffy, J. (2004). "Agent-Based Models and Human Subject Experiments," to appear in K. L. Judd and L. Tesfatsion (eds.) Handbook of Computational Economics vol. 2 Amsterdam: Elsevier.

Easley, D. and J. O. Ledyard (1993). "Theories of Price Formation and Exchange in Double Oral Auctions." in D. Friedman and J. Rust (eds.) *The Double Auction Market: Institutions, Theories, and Evidence* pp. 63–98.

Farmer, J. D., Patelli, P. and Zovko, I. I. (2003). "The predictive power of zero intelligence in financial markets." Preprint, http://xxx.arxiv.org/abs/cond-mat/0309233

Fama, E. F. (1965). "The Behavior of Stock-market Price." *Journal of Business* 39, 226–241.

Gode, D. K. and S. Sunder (1993). "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality." *Journal of Political Economy* 101, 119–137.

Hayek, F. A. (1945). "The use of knowledge in society." *American Economic Review* 35(4), 519–630.

Jamison, J. C. and Charles R. Plott (1997). "Costly Offers and the Equilibration Properties of the Multiple Unit Double Auction Under Conditions of Unpredictable Shifts of Demand and Supply." *Journal of Economic Behavior and Organization* 32, 591–612.

Lawrence, M. (1986). "Weak-form efficiency in the Kuala Lumpur and Singapore stock markets." *Journal of Banking Finance* 10, 431–445.

Perron, P. (1988). "Trends and Random Walks in Macroeconomic Time Series." *Journal of Economic Dynamics and Control* 12, 297–332.

Plott, Charles R. (1982). "Industrial Organization Theory and Experimental Economics." *Journal of Economic Literature* 20, 1485–1587.

Plott, Charles R. and G. George (1992). "Marshallian vs. Walrasian Stability in an Experimental Market." *The Economic Journal* 102, 437–460.

Plott, Charles R. (2000). "Markets as Information Gathering Tools." *Southern Economic Journal* 67(1): 1–15.

Plott, Charles R. (2001). *Markets Institutions and Price Discovery: Collected Papers on the Experimental Foundations of Economics and Political Science, Volume Two* Northhampton, MA: Edward Elgar Publishers.

Smith, V. L. (1962). "An Experimental Study of Competitive Market Behavior." *Journal of Political Economy* 70, 111–137.

Smith, V. L. (1982). "Microeconomic Systems as an Experimental Science." *American Economic Review* 72, 923–955.

Solnik, B. (1973). "A note on the validity of the random walk for European prices." *Journal of Finance* 28, 1151–1159.

Sunder, S. (1995). "Experimental Asset Markets: A Survey" in J. H. Kagel and A. E. Roth (eds.) *The Handbook of Experimental Economics* Princeton University Press, pp. 445–495.

Chapter 3

CHOOSING A MODEL OUT OF MANY POSSIBLE ALTERNATIVES: EMISSIONS TRADING AS AN EXAMPLE

Tatsuyoshi Saijo
Osaka University

Abstract

The main purpose of this paper is to consider how to choose a model when there are many possible alternatives to choose from. We use global warming, especially, emissions trading, as an example. First, we describe each model in a very simple setting and then consider implicit and explicit assumptions underlying each model. In other words, we try to identify the environments in which the model really works. Our models yield results that may be different or occasionally inconsistent. In order to evaluate the results, we argue that the setting of the models and the implications of their implicit assumptions are important.

1. INTRODUCTION

Sulfur dioxide emissions in the atmosphere have detrimental effects on human health through acid rain and soil pollution. Carbon dioxide emissions do not have direct deleterious effects on humans, but they may cause global warming in the future. Because of both the direct and indirect effects of these greenhouse gases (GHGs), a few methods have been proposed to control their emissions into the atmosphere. The traditional method is through direct regulations or the *command and control* method. Another method is *emissions trading* whereby emissions targets are set and agents are given incentives to reduce emissions further since they can sell any surplus emissions permits in the market. In the case of direct regulation, once an agent satisfies the regulation, there is no incentive to reduce emissions further.

The December 1997 Kyoto Protocol to the United Nations Convention on Climate Change called for Annex B countries (i.e., advanced countries and some countries that are in transition to market economies) to reduce their average greenhouse gas emissions over the years 2008–2012 to at least five percent below the 1990 levels. In order to implement this goal, the protocol authorizes three major mechanisms collectively called the *Kyoto Mechanism*. These are 1) emissions trading, 2) joint

A. Rapoport and R. Zwick (eds.), Experimental Business Research, Vol. II, 47–81.

implementation, and 3) the Clean Development Mechanism. As almost no directions are given in the Protocol for implementing these mechanisms, the details of the implementation must be designed.

This paper is organized as follows. In the first part of this paper (sections 2, 3, and 4), we describe models to implement the Kyoto Mechanism by using a marginal abatement cost curve for each country in order to limit the production of greenhouse gases. Because the total amount of greenhouse gases varies over time, dynamic models are required. However, we restrict ourselves to static models here because the aimed period of the Kyoto Protocol is from 2008 to 2012.[1] It is often said that emissions trading attains a fixed goal of regulated emissions at minimum cost. We focus on this statement in Section 2 and show that the emissions reduction cost is minimized at a competitive equilibrium. We then investigate some neutrality propositions. Section 3 introduces a social choice model to consider if competitive equilibrium can be attained through the concept of strategy-proofness. Strategy-proofness means that the best strategy of each country is to report the true marginal abatement cost curve. We will show that a country can gain by not announcing its true marginal abatement cost curve. That is, in the announcement game of marginal abatement cost curves, it is impossible to attain the Kyoto target at the cheapest cost under strategy-proofness. Section 4 proposes a game in which prices and quantities are strategic variables. The possibility of attaining competitive equilibrium through the constructed game is then considered. One such game is the Mitani mechanism (1998), which implements the competitive equilibrium allocation in subgame perfect equilibrium. In GHG emissions trading starting from 2008, the problem of market power is an important issue. Countries such as Russia and Japan will dominate the market. The Mitani mechanism attains the competitive equilibrium allocation even when the number of participants is small. Sections 5 and 6 describe experiments to test the proposed models. Section 5 tests whether the model proposed in Section 2 works in a laboratory setting. We implicitly assume that the price is determined where the demand is equal to the supply, but we cannot determine the prices and quantities without having a concrete trading rule in the laboratory. An important issue is what trading rule should be chosen. Section 6 describes an experiment by Mitani, Saijo and Hamaguchi (1998) that uses the Mitani mechanism. The subjects in this experiment did not choose the subgame perfect equilibrium strategy, but rather cooperated with one another to attain a Pareto superior outcome. We consider the tension between theory and these experimental results and consider why the mechanism did not work well. Finally, Section 7 provides concluding remarks.

2. A SIMPLE MICROECONOMIC APPROACH

Suppose that n countries are involved in emissions trading. Let $MAC_i(x_i)$ be a marginal abatement cost function of country i, where x_i is the amount of green house emissions of i. Suppose, further, that the assigned amount of country i determined by the Kyoto Protocol is $\hat{x}_i$, and the amount of emissions before trading is $\bar{x}_i$. We assume that each country treats the price of emissions permits parametrically. In

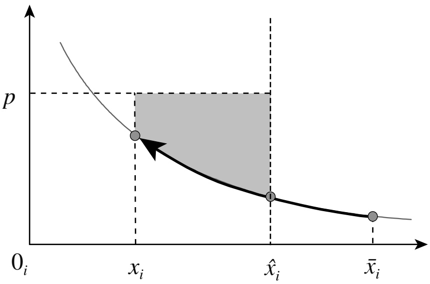

Figure 1-1: A Supplier

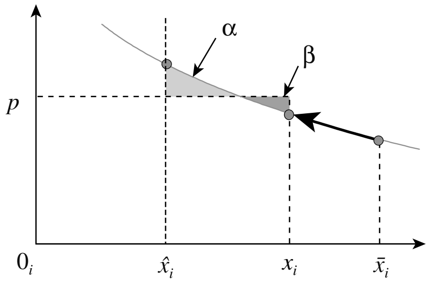

Figure 1-2: A Demander

Figure 1. Determination of GHG Emissions

what follows, we define the profit of each country by the surplus earned relative to the cost of achieving the assigned amount.

Consider first the supplying country in Figure 1-1. The horizontal axis shows the amount of emissions of GHGs, and the vertical axis gives the marginal cost. Country i must reduce its emissions of GHGs from $\bar{x}_i$ to $\hat{x}_i$ in order to attain the target of the Kyoto Protocol. On the other hand, this country could obtain profits from the emissions permit price, p, by selling permits after reducing its GHGs by more than required in the Kyoto target. This country should consider this fact when deciding on the amount of allowable GHG emissions, x_i, in that country. First, the country's total revenue becomes $p(\hat{x}_i - x_i)$. Although the real cost is the area between the marginal abatement cost curve and the line segment, $\bar{x}_i - x_i$, we define the profit function as

$$\pi_i(x_i) = p(\hat{x}_i - x_i) - \int_{x_i}^{\hat{x}_i} MAC_i(t)\,dt, \tag{1}$$

since we consider the profit after attaining the Kyoto target. The shadowed area in Figure 1-1 shows $\pi_i(x_i)$.

Next, consider the demanding country for emissions permits. As Figure 1-2 shows, the domestic reduction cost to attain the Kyoto target is the area between the marginal abatement cost curve and the segment $\bar{x}_i - \hat{x}_i$. On the other hand, reducing emissions by $\bar{x}_i - x_i$ and then buying emissions permits for $x_i - \hat{x}_i$ under the price p makes the payoff

$$\int_{\hat{x}_i}^{\bar{x}_i} MAC_i(t)\,dt - \int_{x_i}^{\bar{x}_i} MAC_i(t)\,dt - p(x_i - \hat{x}_i),$$

which coincides with (1). The payoff (or profit) corresponds to the area $\alpha - \beta$. Notice that the profits or surplus of Figures 1-1 and 1-2 are not maximized under p.

In what follows, we assume that each country maximizes its surplus. Differentiating (1) with respect to x_i, and then finding the maximum, x_i^*, we have the first-order condition

$$p = MAC_i(x_i^*). \tag{2}$$

Furthermore, the sum of emissions for each country must equal the sum of the assigned amount for each country. Thus, we have

$$\sum x_i^* = \sum \hat{x}_i. \tag{3}$$

Therefore, $(p^*, (x_1^*, x_2^*, \ldots, x_n^*))$ satisfying (2) and (3) is a *competitive equilibrium*. Since the first-order condition (2) does not depend on the initial position of emissions, $\bar{x}_i$, the competitive equilibrium does not depend on the initial position or the present amount of emissions from each country. Furthermore, the competitive equilibrium is independent of the distribution of the assigned amounts for each country as long as the total sum of the assigned amounts, $\sum \hat{x}_i$, which is the right-hand side of (3), is fixed. Of course, the initial position and the assigned amount do influence how much cost each country must bear.

Figure 2 illustrates a competitive equilibrium in emissions trading. By superimposing the dotted vertical axis of Figure 1-2 on the dotted vertical axis of Figure 1-1, we create Figure 2 in which the vertical axis at the assigned amount corresponds to these dotted vertical axes. Country 1 reduces its GHG emissions more than the assigned amount and then sells the excess amount over the assigned amount to Country 2. On the other hand, Country 2 reduces emissions by up to

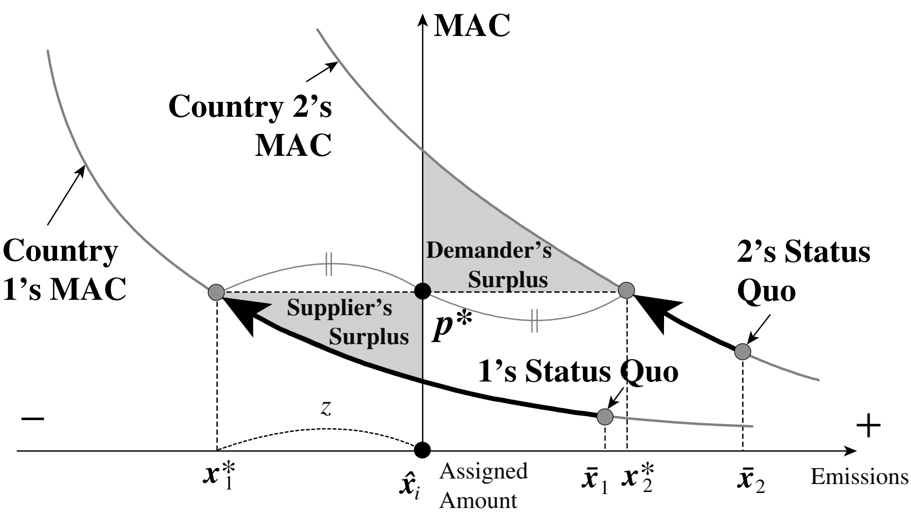

Figure 2. Emissions Trading and the Competitive Equilibrium

x_2^* and then buys emissions permits from Country 1. The marginal abatement cost of each country equals to the competitive equilibrium price.

Consider now a social welfare problem in which the total sum of the reduction costs of all countries is minimized under the condition where the sum of emissions of all countries is equal to the sum of assigned amounts of all countries:

$$\min_x \sum_1^n \int_{x_i}^{\bar{x}_i} MAC_i(t)\, dt \text{ subject to } \sum x_i = \sum \hat{x}_i,$$

where $x = (x_1 x_2, \ldots, x_n)$. If $x^* = (x_1^*, x_2^*, \ldots, x_n^*)$ is a solution to this problem, we have $MAC_i(x_i^*) = MAC_j(x_j^*)$ for all i and j. Consider next a problem where the sum of profits of all countries is maximized under the same condition. Then, the sum becomes

$$\sum_1^n p(\hat{x}_i - x_i) - \min_x \sum_1^n \int_{x_i}^{\hat{x}_i} MAC_i(t)\, dt,$$

since $\sum x_i = \sum \hat{x}_i$, $\sum_1^n p(\hat{x}_i - x_i)$ must be zero for any price p. Therefore, the maximization problem of the sum of profits of all countries becomes

$$\max_x - \sum_1^n \int_{x_i}^{\hat{x}_i} MAC_i(t)\, dt \text{ subject to } \sum x_i = \sum \hat{x}_i.$$

That is, the profit maximization problem is exactly the same as the cost minimization problem. Hence, the solution must satisfy $MAC_i(x_i^*) = MAC_j(x_j^*)$ for all i and j. This also implies that the sum of surpluses (or profits) of all countries is maximized as Figure 2 shows. Summarizing the above results, we have

Proposition 1:
(1) The price of an emissions permit is equal to the marginal abatement cost of each country ($p^* = MAC_i(x_i^*)$ for all i), and the sum of emissions of all countries is equal to the sum of the assigned amounts of all countries at a competitive equilibrium ($\sum x_i^* = \sum \hat{x}_i$).
(2) By minimizing the sum of reduction costs of GHGs of all countries under $\sum x_i = \sum \hat{x}_i$, we have $MAC_i(x_i^*) = MAC_j(x_j^*)$ for all i and j.
(3) The sum of the reduction costs for GHGs of all countries is minimized under a competitive equilibrium.
(4) The sum of surpluses of emissions trading is maximized under a competitive equilibrium.
(5) The competitive equilibrium price and allocation are independent of the initial amount of GHG emissions.

(6) The competitive equilibrium price and allocation are independent of the distribution of the assigned amounts as long as the sum of the assigned amounts of all countries is fixed.

One of the most controversial issues in emissions trading is the so-called "hot air." Although, for example, Russia's assigned amount is 100% of its GHG emissions in 1990, it is estimated that the amount of emissions would be considerably less than the assigned amount after 2008 because of economic stagnation. This implies that Russia could sell emissions permits without actually reducing its GHGs domestically and that buyers of the permits could in turn increase their emissions. This amount of emissions is called "hot air." In Figure 1, $\hat{x}_i - \bar{x}_i$ is hot air as long as $\bar{x}_i$ is on the left-hand side of $\hat{x}_i$. This hot air issue cannot be solved by using emissions trading since the competitive equilibrium allocation is determined independently of hot air as Proposition 1-(5) shows. In other words, the sum of emissions of all countries must be equal to the sum of the assigned amounts of all countries as long as we use emissions trading[3].

The above is a prototype of emissions trading from the viewpoint of microeconomics. We implicitly assumed that the emissions permit can be treated as a private good. However, greenhouse gases are public goods (or "bads"). Behind this approach, we further assume that we can analyze public goods as private goods by introducing a market for emissions permits. But the characteristics of public goods, i.e., non-rivalness and non-excludability, would not disappear with the emissions permit market. In the case of a private good, a transaction is completed if a seller delivers the good to a buyer and the buyer pays for it. On the other hand, in the case of a emissions permit, a seller must reduce GHGs at least by the amount sold and a buyer can emit GHGs by at most the amount bought in addition to the usual transaction of a private good. That is, GHG emissions and reduction must be monitored. It is claimed that actual monitoring of GHG emissions is not an easy task. Furthermore, a time lag exists between the actual emissions and the data collection of the monitoring system even when assuming that the monitoring is perfect. Between these two times, the price of an emission permit will fluctuate in the market. If monitoring and its verification are considerably costly, we cannot be certain that the market for emissions trading minimizes the total costs of emissions reduction.

I did not describe the case when a country emits more than the assigned amount. If this were the case, the violation of the Kyoto target would benefit that country and, at the same time, reduce the country's political trust. That is, that country should emit as much as possible to maximize its economic benefit. In other words, the above model implicitly assumes that there exists some penalty for non-compliance. That is, all countries participating in emissions trading must agree upon a penalty system including the levying of penalties and the form of the organization that levies the penalties. The penalty system necessitates additional costs.

We implicitly assumed that each country treats prices parametrically. However, it is expected that the quantities demanded by Japan and the quantities supplied by

Russia will be relatively large. If they exercise their market power, it is possible that the total surplus will not to be maximized. We will reconsider this problem in Section 5.

The model described in this section is *static*. Further, we assume that each country can move on the marginal abatement cost curve freely. However, it usually takes quite a long time to replace current facilities and equipment. Moreover, the emissions reduction investment might be irreversible. That is, one can move from right to left on the marginal abatement cost curve, but not in the opposite direction. Therefore, marginal abatement costs might not be equalized.

3. A SOCIAL CHOICE APPROACH: STRATEGY-PROOFNESS

Consider a mechanism that determines an allocation through gathering information that each agent has. Under this mechanism, country i reports its own marginal abatement cost curve, MAC_i, to an administrative institution of the United Nations. The institution recommends a competitive equilibrium allocation based upon the reported marginal abatement cost curves. Let $f_i(MAC_1, \ldots, MAC_n)$ be country i's *surplus* accruing from the competitive equilibrium allocation when country i announces MAC_i, and let $f = (f_1, \ldots, f_n)$, be a *surplus function.*

Each country may not necessarily have an incentive to reveal its true marginal abatement cost curve, but the request that each country announces the true marginal abatement cost curve is called strategy-proofness. Let MAC_i^* be country i's true marginal abatement cost curve. Then, a surplus function, $f = (f_1, \ldots, f_n)$, satisfies *strategy-proofness* if, for each country, I, it satisfies

$$(*) \quad \begin{array}{ll} f_i(MAC_1, \ldots, MAC_{i-1}, MAC_i^*, MAC_{i+1}, \ldots, MAC_n) & \text{for all } MAC_j \\ \geq f_i(MAC_1, \ldots, MAC_{i-1}, MAC_i, MAC_{i+1}, \ldots, MAC_n) & (j = 1, \ldots, n). \end{array}$$

That is, strategy-proofness means that announcing the true marginal abatement cost curve is the dominant strategy for all i.

Strategy-proofness is a strong condition since it requires that revealing the true marginal abatement cost curve is the best strategy regardless of the revelations of the others. On the other hand, it would be natural to consider that country i would change its strategy depending on the strategies of the others. This is in the spirit of the Nash equilibrium. Therefore, requiring that the revelation of the true marginal cost curve constitutes a Nash equilibrium, we have

$$(**) \quad \begin{array}{l} f_i(MAC_1^*, \ldots, MAC_{i-1}^*, MAC_i^*, MAC_{i+1}^*, \ldots, MAC_n^*) \\ \geq f_i(MAC_1^*, \ldots, MAC_{i-1}^*, MAC_i, MAC_{i+1}^*, \ldots, MAC_n^*) \end{array} \quad \text{for all } MAC_i$$

for each i. It is clear that the revelation of true marginal cost curve is a Nash equilibrium if the surplus function satisfies strategy-proofness. It may seem that condition (**) is weaker than (*), but these two conditions are in fact equivalent.

Proposition 2 (Dasgupta, Hammond and Maskin, 1979): A surplus function satisfies strategy-proofness if and only if the revelation of the true marginal cost curve for each country, i, is a Nash equilibrium.

We will show that a surplus function does not satisfy strategy-proofness by using Proposition 2. For this purpose, it is sufficient to show that a country could benefit by announcing a false marginal abatement cost curve in a certain profile of true marginal abatement cost curves. Let the solid curves be true marginal abatement cost curves in Figure 3. Assume that country 2 reveals its true marginal abatement cost curve and country 1 announces a false marginal abatement cost curve, shown as a dotted curve. Then, the competitive equilibrium price under this condition becomes P', which is higher than the true competitive equilibrium price. Due to this price increase, the quantity supplied would be less than that of the true competitive equilibrium. Therefore, this reduces the surplus of country 2 by α. On the other hand, country 2 increases its profit by β (= ($P' - P^*$) × the quantity supplied) due to price increase that surpass the loss. Therefore, country 2 could benefit by not revealing its true marginal abatement cost curve. Hence, we have the following proposition.

Proposition 3: A surplus function does not satisfy strategy-proofness.

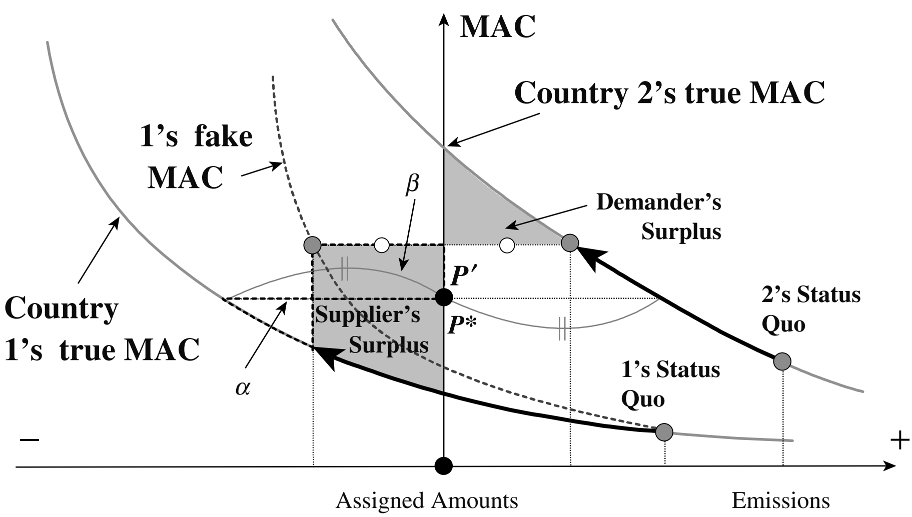

Figure 3. A surplus function is manipulable

Johansen (1977) criticized the research program that considers resource allocation through revelation of functions. Among other things, he pointed out that public goods have never been provided through revelation of utility functions and production functions. Rather, he argued that they have been provided through political processes such as representative systems. Therefore, he indicated that research on public good provision in democratic societies was necessary. Although he expressed this opinion more than twenty years ago, very little progress has been made since then.

There are many ways to evaluate Proposition 3, but we interpret this proposition as saying that the surplus function is manipulable even though transmission of marginal abatement cost curves can be done at a nominal cost through technological innovation. Additional evaluation follows in the next section.

4. MECHANISM DESIGN APPROACH

We consider an emissions allocation problem through the revelation of marginal abatement cost curves assuming that each country does not know the marginal abatement cost curves of the other countries. However, once the emissions trading starts, each country roughly knows the marginal abatement cost curves of the other countries. As Proposition 2 indicates, there is no incentive for each country to reveal its true marginal abatement cost curve since strategy-proofness is equivalent to the condition in which the revelation of true marginal abatement cost curves is a Nash equilibrium. Therefore, apart from the revelation of curves, which requires quite a lot of information, it is worthwhile to consider an allocation problem in which the transmission of prices and quantities of emissions permits is allowed to attain the competitive equilibrium. In what follows, we show that a two-stage mechanism called the Mitani mechanism (1998) attains the competitive equilibrium allocation through information transmission of prices and quantities.[4] The Mitani mechanism works if the number of participants is at least two. That is, the mechanism can overcome the problem of a small number of participants. As stated in Section 2, a country may exercise its market power when the number of countries is small. Even under this condition, the Mitani mechanism attains a competitive equilibrium allocation.

Consider the Mitani mechanism with two countries, 1 and 2. In the first stage, each country announces a price simultaneously. The dotted ellipse in Figure 4 shows that country 2 does not know the price announcement of country 1. Country 2 must

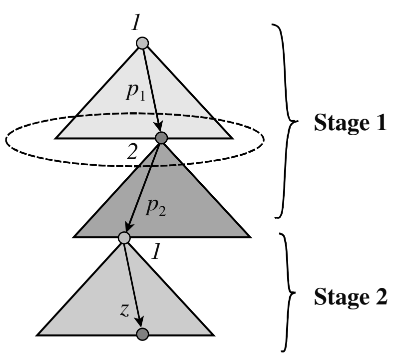

Figure 4. The Mitani mechanism

evaluate its transaction of emissions permits from the announced price, p_1, of country 1, and vice versa. In the second stage, knowing p_1 and p_2 in the first stage, country 1 announces the quantity, z, of emissions permits.

We define the payoff functions of countries 1 and 2 in the Mitani mechanism. For simplicity, we suppose that country 1 is a supplier of emissions permits and country 2 is a demander. Following Section 2, we define the payoff functions in comparison with the case when each country attains the assigned amount domestically. Let z be the quantity supplied by country 1. Then, the cost of reducing from $\hat{x}_1$ by z is

$$C_1(z) = \int_{\hat{x}_1 - z}^{\hat{x}_1} MAC_1(t)\, dt. \tag{4}$$

Since $\hat{x}_1$ is the assigned amount of country 1, the cost for country 1 is a function of z. On the other hand, since the quantity supplied that is determined by country 1 becomes the quantity demanded by country 2 in the Mitani mechanism, the gross payoff of country 2 is

$$C_2(z) = \int_{\hat{x}_2}^{\hat{x}_2 + z} MAC_2(x)\, dx. \tag{5}$$

Following the above, the payoff function, g_i, becomes

$$\begin{aligned} &g_1(z, p_1, p_2) = -C_1(z) + p_2 z - k(p_1 - p_2)^2 \quad \text{where } k > 0 \\ &g_2(z, p_1) = C_2(z) - p_1 z \end{aligned} \tag{6}$$

Country 1 sells its excess reduction, z, beyond its assigned amount to country 2. The price of this transaction is the price that is announced by country 2 in stage 1. The revenue from this transaction is $p_2 z$, and the last term, $k(p_1 - p_2)^2$, in g_1 is a penalty term to make country 1 announce the same price announced by country 2. Country 2 buys z with p_1 announced by country 1.

Consider now how the Mitani mechanism works. First, consider the second stage. Country 1 determines z to maximize the payoff. Therefore, differentiating g_1 with respect to z and equating the result to zero, we have $p_2 = C_1'(z)$. That is, z is determined depending on p_2, which is determined in the first stage. Therefore, we can regard z as a function of p_2. That is, $z = z(p_2)$. By substituting this expression with $p_2 = C_1'(z)$ and differentiating with respect to p_2, we have $z'(p_2) = 1/C_1''$. As long as $C_1'' \neq 0$, $z'(p_2)$ is not zero.

Consider now how country 1 behaves. Country 1 determines p_1 to maximize the payoff. That is, by differentiating g_1 with respect to p_1, we have $p_1 = p_2$. On the other hand, country 2 chooses p_2 to maximize the payoff taking into account the behavior, $z(p_2)$, of country 1. That is, by substituting $z(p_2)$ into g_2 and differentiating with respect to p_2, we have $C_2' \cdot z'(p_2) - p_1 z'(p_2) = z'(p_2)(C_2' - p_1) = 0$. Since $z'(p_2)$ is not zero, we have $p_1 = C_2'(z)$.

Therefore, we have $p_1 = p_2 = C_1'(z) = C_2'(z)$. From (4), we have $C_1'(z) = MAC_1(\hat{x}_1 - z)$ and, from (5), we have $C_2'(z) = MAC_2(\hat{x}_2 + z)$. That is, we have $p_1 = p_2 = MAC_1(\hat{x}_1 - z) = MAC_2'(\hat{x}_2 + z)$. This is exactly the same expression as Proposition 1-(1). This proves that the Mitani mechanism achieves the competitive equilibrium allocation in subgame perfect equilibrium.

The Mitani mechanism consists of a payoff function and a game tree as depicted in Figure 4. If marginal abatement cost curves, MAC_1 and MAC_2, are given, we have a subgame perfect equilibrium (p_1, p_2, z). The subgame perfect equilibrium of the Mitani mechanism is unique. Let the set be $N^{sp}(MAC_1, MAC_2)$. We write $g \cdot N^{sp}(MAC_1, MAC_2)$ as the evaluation of the set by g. Then, we have

$$f(MAC_1, MAC_2) = g \cdot N^{sp}(MAC_1, MAC_2). \tag{7}$$

That is, the payoff function coincides with the surplus function under the Mitani mechanism. In other words, the Mitani mechanism implements the surplus function, f, in subgame perfect equilibrium if (7) holds for any (MAC_1, MAC_2).

Let us find the set of Nash equilibria of the Mitani mechanism. Since the strategic variables of country 1 are p_1, z, by differentiating g_1 with respect to p_1 and z, we have

$$\frac{\partial g_1(z, p_1, p_2)}{\partial p_1} = -2k(p_1 - p_2) = 0, \quad \frac{\partial g_1(z, p_1, p_2)}{\partial z} = -C_1'(z) + p_2 = 0.$$

On the other hand, although the strategic variable of country 2 is p_2, this variable does not affect the payoff of country 1. That is, the payoff of country 2 does not depend on the announcement of p_2. Therefore, any triple (p_1, p_2, z) satisfying $p_1 = p_2 = C_1'(z)$ is a Nash equilibrium of the Mitani mechanism. By requiring $C_1'(z) = C_2'(z)$ in addition to the condition, we obtain a subgame perfect equilibrium. That is, the set of subgame perfect equilibria is a proper subset of the set of the Nash equilibria. Therefore, the Mitani mechanism cannot implement the surplus function under a Nash equilibrium.

We implicitly assume that there must exist a central body that determines who the participants in the mechanism are, collects information, and conducts resource allocation based upon the information that the mechanism requires. In particular, participants must be all countries in the case of global warming, but some country may not want to participate in and benefit from the reduction of greenhouse gases by other countries.[5] As for the collection of information, a participant may not transmit the information required by the mechanism to the body even though the participant agrees to participate in the mechanism. When participants announce their willingness to participate in the mechanism, the central body must have some power to collect information about the strategies of the participants under the framework of social choice or the mechanism design approach. These two approaches implicitly assume that there exists a body with central power.[6] We consider the validity of the Mitani mechanism in the following section.

5. AN EXPERIMENTAL APPROACH

Even though we have succeeded in constructing a theoretical mechanism, we are not certain that it works. Using it in the real world would be risky since the economic damage from a possible failure would be unbearable. Furthermore, we would not be able to determine whether the failure of the mechanism is due to some flaw in the mechanism or some other outside factors. An alternative way to assess the performance of the mechanism is a laboratory experiment. We construct an experimental model out of the theoretical model. Usually, theory does not indicate the number of agents in the model, the exact shape of the functions used in the model, and what types of information each agent has. In an experimental model, these variables must be specified. Depending on how we assign these parameters, we have many possible cases. Our methodology calls for recruiting subjects and paying them contingent upon their performance. By conducting several experiments and then comparing the results, we can understand the effects of these parameters.

We utilize two experiments conducted by Hizen and Saijo (2001) and Hizen, Kusakawa, Niizawa, and Saijo (2000), whose models of the experiments are based upon the microeconomic approach in Section 2. Section 6 describes an experiment that Mitani, Saijo, and Hamaguchi (1998) designed to assess the performance of the Mitani mechanism.

The main question of the experiments by Hizen and Saijo (2001) and Hizen, Kusakawa, Niizawa, and Saijo (2000) is whether the total surplus is maximized under emissions trading. The emissions trading model in Section 2 implicitly assumes that transactions are conducted under a competitive equilibrium, but the surplus of each country can be different from the surplus at the competitive equilibrium depending on the methods of transaction. In order to avoid problems due to non-compliance to the assigned amounts, Hizen and Saijo (2001) designed their experiment so that the assigned amounts are always satisfied in the course of transactions and compared two trading institutions, namely, the double auction and bilateral trading. On the other hand, Hizen, Kusakawa, Niizawa, and Saijo (2000) explicitly incorporated non-compliance and decision making on domestic reductions into their experiment.

Prior to Hizen and Saijo (2001) and Hizen, Kusakawa, Niizawa, and Saijo (2000), Bohm (1997) conducted an important emissions trading experiment. He recruited bureaucrats from the Ministry of Energy and specialists from Finland, Denmark, Norway, and Sweden, and then conducted an emissions trading experiment in which each country could buy *and* sell emissions permits under bilateral trading. The subjects knew the marginal abatement cost curves of other countries, but not the true curves. It took four days to complete a single period by using facsimile communication to exchange information on prices and quantities. The average transaction

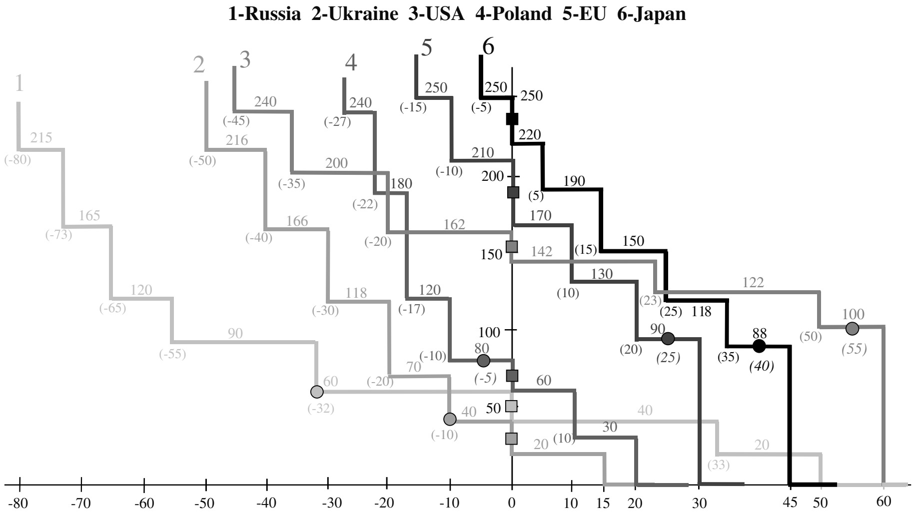

Figure 5. Marginal Abatement Cost Curves used by Hizen and Saijo and Hizen, Kusakawa, Niizawa, and Saijo

price was very close to the competitive equilibrium price and the efficiency was 97%, which is quite high. Many economists have expressed the opinion that it is difficult to attain efficiency if a trading agent is a country as a unit, but Bohm has shown that it is possible to attain high efficiency when countries are the players.

I start with an overview of the common features of the experiments of Hizen and Saijo (2001) and Hizen, Kusakawa, Niizawa, and Saijo (2000). Six subjects participated in a session. The subjects were supposed to represent Russia, Ukraine, the US., Poland, the EU, and Japan. In the experiments, no country names were given to subjects. Subjects must have assumed that they were engaged in transactions of an abstract commodity. Figure 5 shows the marginal abatement cost curves used in the experiments. The origin is the assigned amount for each country. In the experiments on information disclosure on the marginal abatement cost curves, every subject had Figure 5 at hand. On the other hand, in the experiments on information concealment, each subject only knew his/her marginal abatement cost curve.

Two trading methods were used in the experiments. The first one was bilateral trading. Two out of six subjects made a pair and then negotiated the price and quantity of the emissions permits. The maximum number of pairs was three because of the limit of six subjects. During the negotiation, subjects were not allowed voice communication, but communicated by means of writing the numerical values of price and quantity. Written responses of "yes" and "no" were allowed. Once a pair reached an agreement, the pair was supposed to inform the experimenter. In the case of disclosure of information of the contract, the experimenter wrote the information on the blackboard and announced it to every subject. The pair again returned to the floor to seek other contracts. This procedure lasted up to 60 minutes.

Table 1. An Example of the Double Auction

Buyers' Bids	*Sellers' Asks*
(3) \$56, 20 units (1) \$86, 13 units (2) grabs (4)′ ask	(6) \$104, 15 units (4) \$92, 20 units
⋮	⋮

The second method was a double auction. Six subjects were placed together. Subjects who wanted to sell announced the price and quantity, and subjects who wanted to buy made similar announcements. Table 1 shows an example of the auction. A subject who wanted to announce price and quantity raised his/her hand. The experimenter called on the subject (subject 3 in the example since he/she was the first person to raise his/her hand). The subject then called out that he/she wants to buy 20 units at \$56 for each unit, and the experimenter wrote the information on the blackboard. Right after this announcement, another subject (subject 6 in the example) announced his/her willingness to sell 15 units at \$104 for each unit. Since the spread of the price difference was quite high, subject 1 announced his/her willingness to buy 13 units at \$86. Then, subject 4 announced his willingness to sell 20 units at \$92. Right after this announcement, subject 2 accepted the offer from subject 4. The maximum number of units was 20. This process then continued. An important feature of the double auction is that each subject receives the information of the announcements simultaneously. In contrast, in the case of bilateral trading only a pair knows the information as long as the experimenter does not reveal it.

Next, we consider the Hizen and Saijo's (2001) experiment. In order to avoid the non-compliance problem, the starting point of the transaction for each subject was the assigned amount (see squares at the vertical axis going through the origin in Figure 6). When this is the case, the target of the Kyoto Protocol is automatically satisfied at any point in the transaction. As Propositions 1-(5) and 1-(6) show, the competitive equilibrium in Figure 2 coincides with the one in Figure 6. The squares in Figure 6 show the initial points of transactions. Furthermore, Hizen and Saijo assumed that subjects can move on the marginal cost curves freely in order to avoid investment irreversibility in which a subject cannot go back to the left once he/she decides to choose a point on the curve.

Subjects were paid in proportion to their performance in the experiment. In the Hizen and Saijo's (2001) experiment, subjects were instructed about how they could obtain monetary reward by showing them a sample graph such as Figure 7. The point of departure was the assigned amount in Figure 6 that corresponded to "your position" in Figure 7. Consider, as an example, the upper central graph in Figure 6. The horizontal line shows the price line. Since the marginal abatement cost is higher than the emissions permit price, the subject can benefit by buying a permit. If he/she succeeds in buying the permit up to the intersection point between the

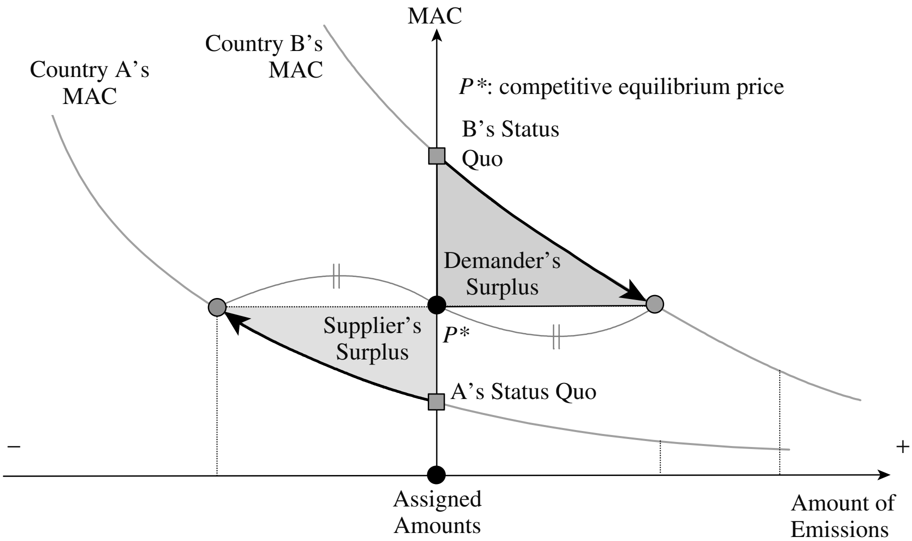

Figure 6. Initial Points of Transactions and Emissions Trading in Hizen and Saijo's Experiment

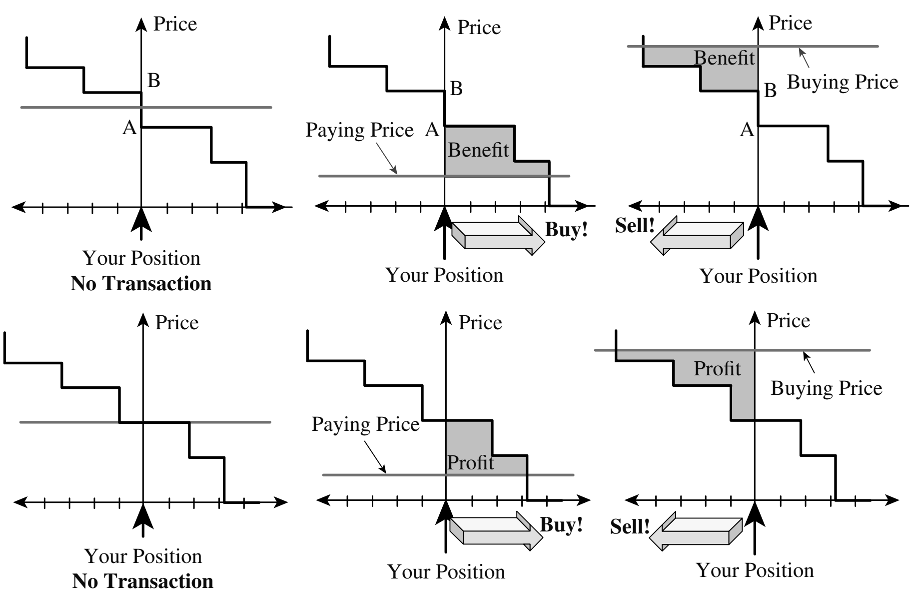

Figure 7. Benefits and Profits of Subjects

Table 2. Controls in Hizen and Saijo's Experiment

			Marginal Abatement Cost Curve Information	
			Disclosure (O)	Closure (X)
Bilateral Trading	Contract Information	Disclosure (O)	**OO**	**OX**
		Closure (X)	**XO**	**XX**
Double Auction			Disclosure (O)	Closure (X)
			O	**X**

marginal cost curve and the price line, he/she could obtain the benefit corresponding to the benefit area in Figure 7. After the transaction, the position of the subject moves. A new transaction starts from this new position. Figure 7 shows all the possible cases.

Consider next the controls in Hizen and Saijo's experiment. In the bilateral trading setting, two controls were used. The first was the marginal cost curve information depending on if the subjects knew all the marginal cost curves. The second was the contract information depending on the concealment or disclosure of the contracted price. Although the assumption that each country knows all the marginal abatement cost curves is unrealistic, we employed it because Bohm (1997) observed high efficiency under the condition that every country knows the curves fairly well. By comparing the results with full information of marginal abatement cost curves to the ones with only private marginal abatement cost information, we can measure the effect of the information disclosure. In the case of the double auction, only the marginal abatement cost information control is employed since the contract information is available to every country. For each cell in Table 2, at least two sessions were conducted with different sets of subjects.

Consider Hizen and Saijo's (2001) experimental results from the viewpoint of efficiency. As Proposition 1-(4) shows, the total surplus is maximized at the competitive equilibrium. That is, the total surplus accruing from any trading rule cannot exceed the total surplus of the competitive equilibrium. Therefore, define a measure of efficiency as follows:

$$\frac{\textit{The sum of the surplus of all subjects}}{\textit{The total surplus of the competitive equilibrium}}$$

Therefore, the maximum efficiency is one or 100%. In the Hizen and Saijo's (2001) experiment, the competitive equilibrium price is between 118 and 120. Therefore, we take the average 119 as the competitive equilibrium price. When the transaction is conducted under this price, the total sum of the surplus is 6990. The first row in Table 3 shows the trading rule, and the second shows the names of the sessions. For example, “OX2” under bilateral trading indicates the disclosure of contract prices, the closure of marginal cost curve information, and the second session of this condition. That is, “O” indicates “disclosure” and “X” indicates “concealment.” The first digit indicates the contract information; the second one indicates the marginal cost curve information; and the last digit indicates the session number. Since the contract information is disclosed in the double auction, “X1” indicates the concealment of the marginal cost curve information and the first session of this control. The number in the leftmost column indicates the subject number; the number in parentheses indicates the surplus obtained by the subject at the competitive equilibrium. In each cell, the upper number is the surplus that the subject actually obtained in the experiment, while the lower number is the individual efficiency. For example, the individual efficiency of subject 1 is 0.732, which is obtained by the ratio of the surplus at the competitive equilibrium (2555) to the actual surplus (1870). The individual efficiency is different from the session efficiency.

The former may exceed one since the distribution of total surplus depends on how the transactions are carried out although the total sum of the surplus cannot exceed 6990. Some subjects might have low efficiency since they sell their emissions permits for less than the competitive price. On the other hand, some might attain high efficiency by buying permits at low prices. In the experiment, the maximum monetary reward was 7600 yen, the minimum 2000 yen, and the average 3459 yen.

As Table 3 shows, the efficiency of each session was quite high except for the “XO1” session. Subject 5 in this session traded even though he/she suffered considerable losses. On the other hand, individual efficiencies varied even in subjects who had the same i.d. number. In bilateral trading, Russia’s efficiencies were lower than those of the competitive equilibrium, and Poland’s efficiencies were higher than those of the competitive equilibrium. We cannot say that the efficiencies of the other countries were statistically different from one another. On the other hand, under the double auction, the efficiencies of the US were higher than those of the competitive equilibrium, Poland’ efficiencies were higher than those of the competitive equilibrium, and the efficiencies of the others were statistically close to one.

Consider now the efficiency dynamics over time. Under bilateral trading, the efficiencies of all sessions, except for session XX2, exceeded 80%, and the efficiencies immediately after 25 minutes exceeded 90% in six out of eight sessions. The efficiencies increased monotonically, but the efficiency of session OX2 fluctuated since one subject bought permits at a loss and then sold some of them at a relatively high price. Under the double auction, the efficiencies of all sessions, except for session X2, exceeded 70% immediately after 17 minutes and exceeded 90% after 44 minutes. They increased monotonically, but we observed a decrease in the X sessions after achieving 100%. For example, in session X2, one subject sold permits at a loss while expecting to buy them at a relatively low price, but could not. Such losses were not observed when the marginal cost curve information was disclosed. These results led to the following observation:

Table 3. Efficiencies in Hizen and –Saijo's experiment

	Bilateral Trading								*Double Auction*				
Subject No.	*OO1*	*OO2*	*OX1*	*OX2*	*XO1*	*XO2*	*XX1*	*XX2*	*O1*	*O2*	*O3*	*X1*	*X2*
1 (2555) (Russia)	1420 0.556	1870 0.732	960 0.376	1710 0.669	1510 0.591	1100 0.431	1460 0.571	1600 0.626	2410 0.943	2410 0.943	1981 0.775	2260 0.885	2865 1.121
2 (1290) (Ukraine)	1140 0.884	914 0.709	360 0.279	1665 1.291	1320 1.023	940 0.729	1536 1.191	2370 1.837	1320 1.023	1320 1.023	520 0.403	1770 1.372	1120 0.868
3 (610) (U.S.A.)	685 1.123	683 1.120	2060 3.377	372 0.610	1846 3.026	615 1.008	583 0.956	550 0.902	850 1.393	865 1.418	1144 1.875	681 1.116	1270 2.082
4 (390) (Poland)	520 1.333	570 1.462	850 2.179	530 1.359	500 1.282	555 1.423	910 2.333	500 1.282	200 0.513	350 0.897	230 0.590	209 0.536	355 0.910
5 (620) (EU)	800 1.290	1105 1.782	1300 2.097	755 1.218	−150 −0.242	1080 1.742	81 0.131	150 0.242	750 1.210	500 0.806	1380 2.226	700 1.129	0 0.000
6 (1525) (Japan)	2425 1.590	1800 1.180	1450 0.951	1844 1.209	1400 0.918	2700 1.770	2390 1.567	1800 1.180	1430 0.938	1515 0.993	1695 1.111	1350 0.885	1360 0.892
Sum (6990)	6990 1	6942 0.993	6980 0.999	6876 0.984	6426 0.919	6990 1	6960 0.996	6970 0.997	6960 0.996	6960 0.996	6950 0.994	6970 0.997	6970 0.997

Observation 1.
(1) *The efficiency of bilateral trading is almost one, regardless of concealment or disclosure of price and marginal abatement cost information.*
(2) *Russia's efficiency is low; Poland's efficiency is high; and the efficiency of the other countries is close to one.*
(3) *The efficiency of allocation in 6 of 8 sessions was more than 90% after 25 minutes.*

We employed the following method to check the convergence of the transaction prices. If the variance of the last three transactions was significantly smaller than that of the first three transactions, we said that the transaction price sequence converged. We also regarded the average of the prices of the last three transactions as the converged price if the sequence converged. This resulted in the following observation.

Observation 2 (Bilateral Trading).
(1) *The contracted average prices in the "XX" sessions (concealment of prices and concealment of marginal abatement cost curves) were roughly equal to the competitive equilibrium price, but the variances in the contracted prices in the "XX" sessions were larger than those in the rest of the sessions.*
(2) *The contracted average prices cannot be said to equal the competitive equilibrium price in sessions other than the "XX" sessions.*
(3) *The average price of the last three contracts is not equal to the competitive equilibrium price in any session.*
(4) *Convergence of the contracted prices is found in five of eight sessions, but no information disclosure effect on convergence is observed.*

Under bilateral trading, the price that a pair agrees on is determined by the negotiation. Even when the transaction prices are open to every subject, a proposal such as "let us agree upon the price that is used by other subjects" can be rejected by the other subject. In other words, negotiation power is an important factor, and, hence, it would be difficult to say that the sequence converged to the competitive equilibrium price. On the other hand, in the double auction, the price sequence converged to the competitive equilibrium price, leading to the following observation:

Observation 3 (Double Auction).
(1) *The contracted price sequence converged to the competitive equilibrium price regardless of concealment or disclosure of the price and marginal abatement cost information.*
(2) *The average price in sessions O1, O2 and X1 were close to the competitive equilibrium price.*
(3) *The average price of the last three transactions in sessions O1, X1 and X2 were close to the competitive equilibrium price.*

In order to understand the effect of the disclosure of transaction information, we should compare the results in sessions OO and XO with sessions OX and XX, respectively. However, since we found that the variance in the contracted prices in session OX1 is significantly different from that in session OX2, we excluded them in the comparison. Instead, we compared sessions OO and XO, where no difference in variance was observed. Similarly, the effect of the disclosure of the marginal abatement cost curves can be seen by comparing sessions XO and XX. We found that the variance in sessions XO is statistically smaller than that in sessions XX. Summarizing the above results, we make the following observation:

Observation 4 (Bilateral Trading).
(1) *Assuming that the marginal abatement cost curves are public information, the disclosure of contracted prices does not have any impact on the variance of in contracted prices.*
(2) *Under the concealment of contracted prices, the disclosure of marginal abatement cost curves reduces the variance of contracted prices.*

Let us now consider the effect of disclosure of the marginal cost curve information. The variance in the O sessions is smaller than that in the X sessions. Russia, Ukraine, and Poland are sellers and the US., EU, and Japan are buyers at the competitive equilibrium. We observed that this happened in the O sessions, but there was at least one subject who bought and sold permits. The numbers of transactions were 9, 8 and 7 in sessions O1, O2, and O3, respectively, and they were 11 and 12 in sessions X1 and X2, respectively. Summarizing the above findings, we make the following observation:

Observation 5 (Double Auction). *The disclosure of marginal abatement cost curves:*
(1) *reduces the variance of contracted prices;*
(2) *makes Russia, Ukraine and Poland only sell, and the US, EU and Japan only buy;*
(3) *reduces the number of trades.*

Aside from the efficiencies, we can also see how the marginal abatement costs changed over time. Due to the step-function nature of our marginal abatement cost curve, we must be careful when evaluating marginal costs. For example, the marginal abatement cost of Russia in session "OO1" was 90 after 25 minutes. In the raw data, we find that Russia sold exactly 55 units of emissions allowances in 25 minutes. Therefore, if the subject wanted to sell one more unit, its marginal abatement cost would have been 120 (see Figure 5). Taking account of this fact, we make the following observation:

Observation 6.
(1) *In bilateral trading, except for the EU subject in session "XO1," the marginal abatement costs of all subjects approach the competitive equilibrium price, but the contracted prices do not.*
(2) *The Marginal abatement costs of all subjects approach the competitive equilibrium price.*

In our double auction experiment, marginal abatement costs converged less rapidly than in the bilateral trading setting. We conjecture that this arises because at most one pair in the double auction can trade at the same time while at most three pairs can do so under bilateral trading.

In order to understand how much market power a country has, we need an aggregate excess demand curve of all the subjects regarding the marginal abatement cost curves as the excess demand curves for emissions allowances. In our design, the competitive equilibrium price range is from 118 to 120, while the excess demand for permits is zero under this price range. Each country might be able to change the equilibrium price by increasing (or decreasing) the quantity supplied (or demanded). If so, and the surplus of this country under the new equilibrium price is greater than the surplus under the true equilibrium price. Then we say that the country has *market power*. After careful examination, we find that the only country that has market power in our design is the US. Table 3 shows that the benefits of the US were more than three times larger than the benefit at the competitive equilibrium in two out of eight sessions under bilateral trading. A statistical test shows that the US did not exercise its market power in any session. Most probably, the subjects could not exploit the marginal abatement cost curve information to use their market power. Under double auction, the individual efficiency of the US is statistically greater than one.

It is remarkable to find that high efficiency was observed even when there existed a subject who had market power. What would happen if subjects could easily find out that they have some market power and the transaction is done by double auction? Bohm (2000) found that the efficiency in this setting is still high, but the distribution of the surplus is distorted. That is, it is often said that the efficiency of the market would be damaged when there are countries that have market power, but this is not confirmed in laboratory experiments. It seems that the double auction and the typical explanation of a monopoly are totally different from each other. In a textbook theory of monopoly, a monopolist offers a price to every buyer, and a buyer must accept or reject the price. The second point is that a country that is supposed to be a seller under the competitive equilibrium price would be a buyer if the price of permits were considerably low.

Consider the policy implications of Hizen and Saijo's experiment. If the main target of a policy maker is efficiency in achieving the Kyoto target, both bilateral trading and the double auction can attain this goal. If the policy maker's target is to achieve equity so that the same permits must be traded at the same price, the double auction is better than bilateral trading. If market power is not exercised, then it seems that bilateral trading is better than the double auction. If the policy maker believes that the information transaction takes a considerable amount of resources, then the double auction is better than bilateral trading.

Hizen, Kusakawa, Niizawa and Saijo (2000) focus on two assumptions that are employed by Hizen and Saijo (2001). The first is that the starting point of the transaction in Hizen and Saijo is the assigned amount of the Kyoto target. The second is that a country can move on the marginal abatement cost curve freely. This assumption is made to avoid the non-compliance problem. In Hizen, Kusakawa, Niizawa and Saijo's (2000) experiment, the starting point of the transaction becomes more realistic as a circle on the marginal cost curve shown in Figure 5. Furthermore, they impose two restrictions on the movement on the marginal cost curve. A country can move on it from right to left, but not in the opposite direction. Once a country spends resources for abatement, it cannot reduce its marginal abatement costs through increased emissions. This corresponds to investment irreversibility. Once an agent invests some resources, the agent cannot go back to the original position. The second restriction is a condition on the decision making on domestic abatement. During 60 minutes of transactions, a country must decide on its domestic abatement decision within half an hour. This reflects that it takes a considerable amount of time to reduce emissions after the decision is made. On the other hand, emissions trading can be conducted any time during the 60 minutes. Under these new conditions, a country might not be able to attain the assigned amount of emissions. If this is the case, then the country must pay a penalty of \$300 per unit. This is considerably high since the competitive equilibrium price range is from \$118 to 120.

In Hizen, Kusakawa, Niizawa, and Saijo's (2000) experiment, the marginal abatement cost curves are private information. The trading methods are bilateral trading and double auction. In bilateral trading, the control is the disclosure of contract price information (O) or the concealment of this information (X). In the double auction, this information is automatically revealed to everyone. The rest of the design is the same as in Hizen-Saijo's experiment.

Table 4 is similar to Table 3. Let us explain the two numbers under the name of each country. The US has (55, 50), for example, indicating that the initial point is 55 and the competitive equilibrium point after the transaction is 50 (Figure 5). Now, consider the two numbers in the data. The numbers for the US in session O4 in bilateral trading is (23, –2). The first number shows that the US conducted 32 (= 55 – 23) units of domestic reduction, which resulted in 23 units on the horizontal axis by moving the marginal abatement cost curve. In order to comply with the Kyoto target, the US must buy at least 23 units of emissions permits, but since the US bought 25 units, this resulted in –2 on the horizontal axis. That is, the US achieved 2 units of over-compliance.

We have two kinds of efficiency. The first is the actual efficiency attained. That is, actual efficiency measures the actual surplus attained in each experiment after assigning a zero value to units of over-compliance and \$300 for each unit of non-compliance as \$300. This is shown in the bottom row of Table 5. For example, the actual surplus of session O4 is 5736 and its efficiency is 0.821. The second kind of efficiency is the modified efficiency that reevaluates units of over-compliance and units of non-compliance by using the concept of opportunity costs. Details are given in Hizen, Kusakawa, Niizawa, and Saijo (2000). This is shown underneath the box in Table 5. For example, the modified surplus of session O4 is 6596 and its modified efficiency is 0.944. The average efficiency (the modified efficiency) in the X sessions is 0.605 (0.811); in the O sessions it is 0.502 (0.807); and in the D sessions it is 0.634 (0.873).

After a careful look at Table 5, we make the following observation.

Table 5. Efficiencies in the Hizen, Kusakawa, Niizawa, and –Saijo's (2000) experiment

	Bilateral Trading								*Double Auction*			
Subject No.	*X1*	*X2*	*X3*	*X4*	*O1*	*O2*	*O3*	*O4*	*D1*	*D2*	*D3*	*D4*
1 (2555) (Russia) (–32)(–55, 0)	1535 0.601 –40, 0	1600 0.626 –32, 0	620 0.243 –65, –22	656 0.257 –42, 0	1415 0.554 –55, –3	384 0.150 –52, 0	1825 0.714 –33, 0	1465 0.573 –52, 0	1425 0.558 –55, 0	2435 0.953 –65, 0	1360 0.532 –44, 0	2060 0.806 –60, 0
2 (1290) (Ukraine) (–10)(–30, 0)	766 0.594 –28, 0	1175 0.911 –20, 0	1820 1.411 –30, 0	700 0.543 –20, 0	–565 –0.438 –20, –15	2625 2.035 –20, 0	1285 0.996 –25, 0	2200 1.705 –20, 0	1195 0.926 –30, 0	–30 –0.023 –30, –27	850 0.659 –20, 0	–1925 –1.492 –30, –37
3 (610) (U.S.A.) (55)(50, 0)	1046 1.715 23, 0	220 0.361 *30, 3*	556 0.911 23, 0	1416 2.321 23, 0	–4130 –6.770 –20, –30	–4094 –6.711 *50, 23*	481 0.789 23, 0	316 0.518 23, –2	890 1.459 40, 0	641 1.051 23, 0	769 1.261 23, 0	–404 –0.662 23, 0
4 (390) (Poland) (–5)(–10, 0)	240 0.615 –5, 0	100 0.256 –10, 0	20 0.051 –17, 0	94 0.241 –10, 0	77 0.197 –10, 0	500 1.282 –10, 0	300 0.769 –13, 0	450 1.154 –10, 0	165 0.423 –10, 0	275 0.705 –10, 0	375 0.962 –11, 0	763 1.956 –17, 0
5 (620) (EU) (25)(20, 0)	–650 –1.048 5, –5	375 0.605 10, 0	850 1.371 25, 0	850 1.371 20, 0	1002 1.616 20, 0	975 1.573 20, 0	630 1.016 20, 0	965 1.556 20, –2	760 1.226 20, 0	–900 –1.452 10, –10	770 1.242 20, 0	682 1.100 20, 0
6 (1525) (Japan) (40)(25, 0)	2175 1.426 35, 0	2130 1.397 25, 0	–3100 –2.033 15, –25	1710 0.245 *35, 6*	1931 1.266 35, 0	2040 1.338 25, 0	1625 1.066 25, 0	340 0.223 30, –5	2515 1.649 35, 0	25 0.016 25, –10	1822 1.195 25, –5	1200 0.787 25 , 0
Sum (6990)	5112 0.731	5600 0.801	766 0.110	5426 0.776	–270 –0.039	2430 0.348	6146 0.879	5736 0.821	6950 0.994	2446 0.350	5946 0.851	2376 0.340
	5612 0.803	6230 0.891	4136 0.592	6686 0.957	3140 0.449	6680 0.956		6596 0.944		5856 0.838	6426 0.919	5186 0.742

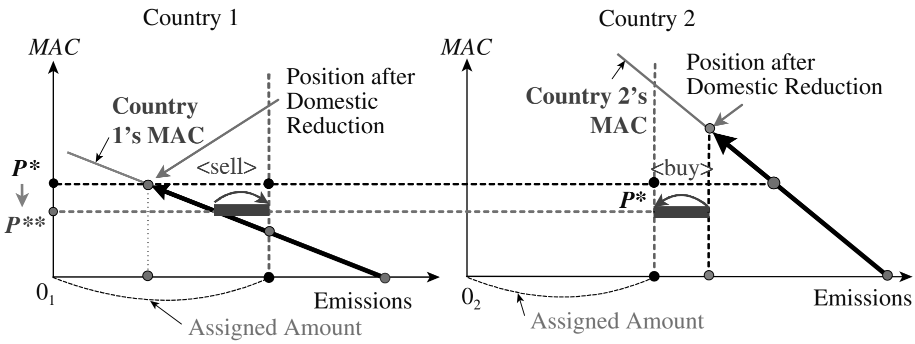

Figure 8. Point Equilibrium

Observation 7.
(1) *Russia's domestic reductions were not enough in bilateral trading, but they were close to the domestic reduction at competitive equilibrium in the double auction.*
(2) *The US conducted excessive domestic reductions in all sessions.*
(3) *In bilateral trading, nine cases of over-compliance and three cases of non-compliance out of 48 cases were observed. On the hand, in the double auction, five cases of over-compliance and no case of non-compliance out of 24 cases were observed.*

In order to understand the nature of investment irreversibility, Hizen, Kusakawa, Niizawa, and Saijo (2000) introduced a point equilibrium. In Figure 8, the competitive equilibrium price is P^*. If country 2 continues to climb the marginal abatement cost curve, the price that equates the quantity demanded and the quantity supplied should go down and it should be P^{**}. We call this "should be" price the point equilibrium price. Even though the point equilibrium price is P^{**}, countries might have been trading permits at a higher price than P^*.

In each session, we have two pieces of price sequence data. One is the actual price, and the other is the point equilibrium price. With the help of the point equilibrium price, we found two types of price dynamics. The first is the early point equilibrium price decrease case (or type 1), and the second is the constant point equilibrium price case (or type 2). We observed five sessions of type 1 and seven sessions of type 2 out of 12 sessions.

Figure 9 shows two graphs of type 1 and type 2 price dynamics. The top picture shows a typical case of type 1 and the bottom a typical case of type 2. The horizontal axis indicates minutes, and the vertical axis prices. The horizontal line indicates the competitive equilibrium price, and the dark step lines indicate the point equilibrium prices. A box indicates a transaction. The left-hand side is the seller's name; the right hand side is the buyer's name; and the bottom number indicates the number of units in the transaction. A diamond indicates the domestic reduction. Consider the top graph that is for session D2. Up until 15 minutes, we observe many diamonds that indicate

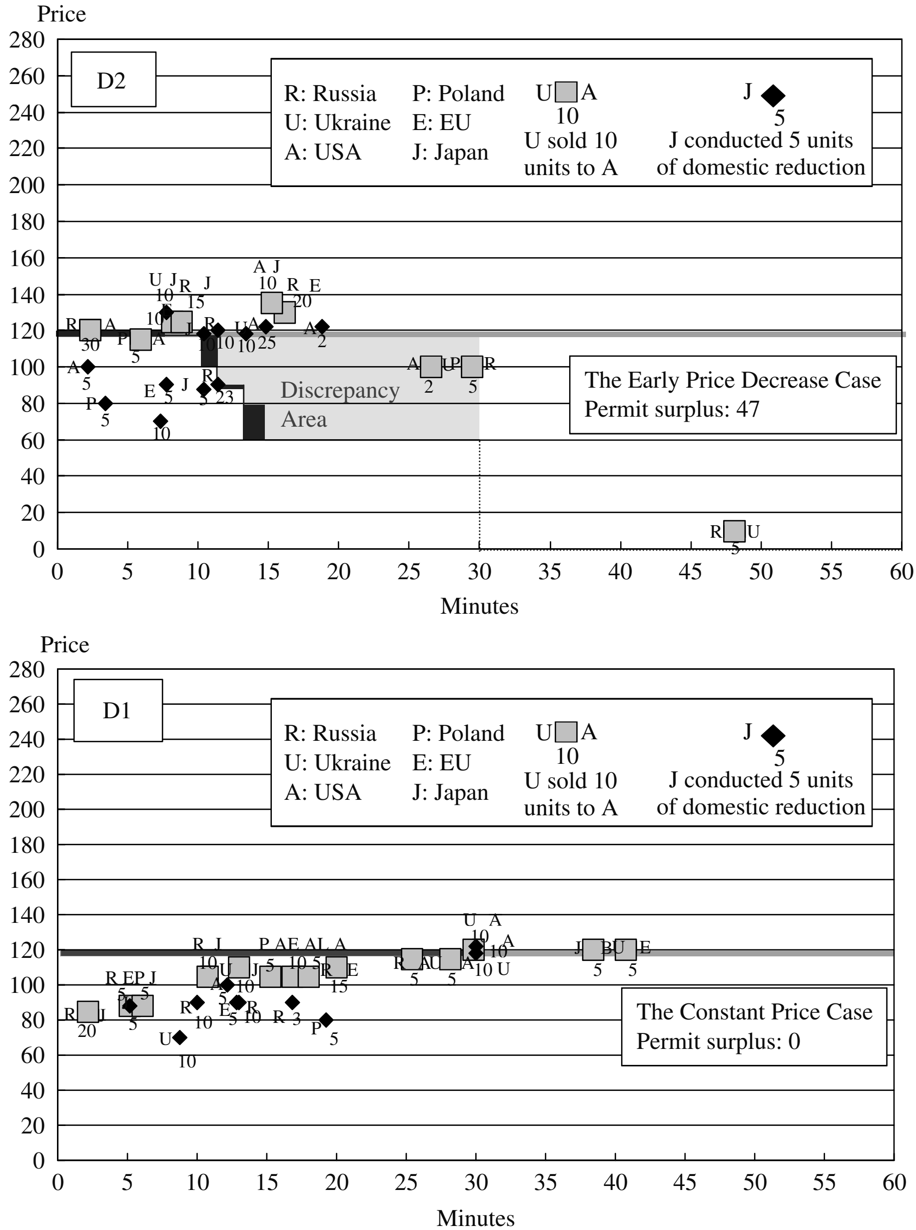

Figure 9. Price Dynamics of Hizen, Kusakawa, Niizawa, and Saijo's experiment

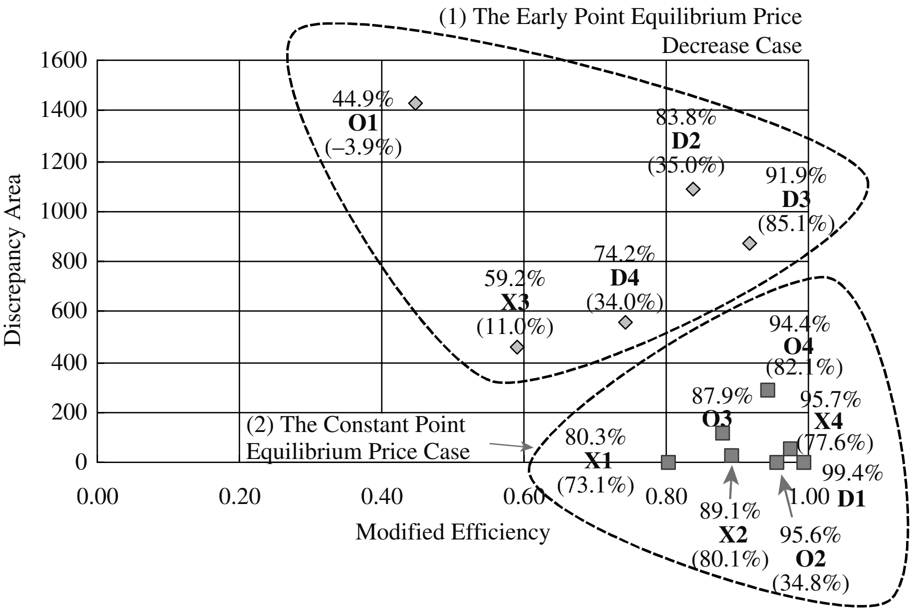

Figure 10. The Relationship between Modified Efficiency and the Discrepancy Area

domestic reduction. This reduction seems to come from fear of non-compliance of demanders. This causes the the transaction price to be higher. Even after the point equilibrium price decreased after 10 minutes, the actual transaction prices were considerably higher than point equilibrium prices. That is, high price inertia was observed. After half an hour, no domestic reduction was possible and the point equilibrium becomes zero. We measured the area between the competitive equilibrium price line and the point equilibrium price curve up to half an hour as the discrepancy area.

In the case of the bottom graph, the starting price was relatively low. Due to this low price, supply countries did not conduct enough domestic reduction. After 10 minutes and until 30 minutes, the demand countries conducted considerable domestic reduction. In this case, the point equilibrium price curve coincided with the competitive equilibrium price line. That is, the discrepancy area was zero.

Figure 10 illustrates the relationship between the modified efficiency and the discrepancy area. By cluster analysis, we have found two groups, type 1 and type 2. Although the number of sessions was quite small, within the same type, it seems that efficiencies of the double auction were higher than those of bilateral trading and that information disclosure increased the efficiency. Summarizing these findings, we make the following observation:

Observation 8.
(1) *Two types, i.e., the early point equilibrium price decrease case and the constant point equilibrium price case were observed.*

(2) *Excessive domestic reduction was observed in both types.*
(3) *In both types, efficiencies in the double auction were higher than those in bilateral trading.*
(4) *In type 1, we observed high price inertia and a sudden price drop.*
(5) *In type 2, insufficient domestic reduction from the supply countries caused excessive domestic reduction from the demand-countries.*

The sudden price drop observed in Observation 8–(4) would be overcome by banking, which is allowed in the Kyoto Protocol. Muller-Mestelman (1998) found that banking of permits had some power to stabilize the price sequence. Furthermore, under either trading rule, early domestic reduction resulted in type 1 and caused a efficiency lower than that of type 2. It seems that haste makes waste.

6. EXPERIMENTAL APPROACH (2)

This section describes the experimental results of Mitani, Saijo, and Hamaguchi (1998) who studied the Mitani mechanism. In their experiment, they specify cost $C_1(z)$ and $C_2(z)$, as follows.

$$C_1(z) = 37.5 + 0.5(5 + z)^2, \quad C_2(z) = 15z - 0.75z^2.$$

Furthermore, the penalty function of country 1 is specified by

$$d(p_1, p_2) = \begin{cases} 0 & \text{if } p_1 = p_2 \\ K & \text{if } p_1 \neq p_2 \end{cases}, \text{ where } K > 0.$$

Thus, if countries 1 and 2 announce the same price, then the penalty is zero; if not, then the fixed amount of penalty is imposed on country 1. Therefore, the payoff functions of the mechanism become

$$g_1(z, p_1, p_2) = -C_1(z) + p_2 z - d(p_1, p_2).$$
$$g_2(z, p_1) = C_2(z) - p_1 z$$

Even with modification of the Mitani mechanism, the subgame perfect equilibrium would not be changed. Applying the condition of subgame perfect equilibrium to the Mitani mechanism, $p_1 = p_2 = C_1'(z) = C_2'(z)$, we have $C_1' = 5 + z$, $C_2' = 15 - 1.5z$, and hence $z = 4$. That is, $p_1 = p_2 = 9$.

The experimental test of the Mitani mechanism is designed so that each agent is supposed to minimize the cost. Therefore, by putting a minus sign in the payoff of country 1, we have

The total cost of country 1 = 37.5 + 0.5 × (5 + the units of transaction)^2 − (the price that country 2 chose) × (the units of transaction) + the charge, where the charge term is $d(p_1, p_2)$. We regard the payoff of country 2 as the surplus accruing from buying emissions permits from x_2^* to the assigned amount ($\hat{x}_2$) as shown in Figure 2. On the other hand, in the Mitani, Saijo, and Hamaguchi's experiment, the sum of the cost-reducing emissions from $\bar{x}_2$ to x_2^* and the payment $p^*(x_2^* - \hat{x}_2)$ for emissions are the total cost. This does not change the subgame perfect equilibrium of the Mitani mechanism since it merely changes the starting point for either the payoff or cost. When $\bar{x}_2 = 10$, $C_2(10) - C_2(z) = 75 - 15z + 0.75z^2 = 0.75(10 - z)^2$ is the cost of reducing the amount of emissions from $\bar{x}_2$ to x_2^*. That is,

The total cost of country 2 = 0.75 × (10 − the units of transaction)^2 + (the price that country 1 chose) × (the units of transaction).

When no transaction occurs, the total cost of country 1 is $37.5 + 0.5 \times 5^2$, which equals 50 and the total cost of country 2 is 0.75×10^2, which is 75, where the charge term is zero.

Let us review the experiment. Two sessions were conducted, one for $K = 10$ and the other for $K = 50$. Each session included 20 subjects who gathered in a classroom and divided into 10 pairs. Each subject could not identify the other dyad member. During the experiment, "emissions trading" were not used. Country 1 in the above corresponded to subject A, and country 2 to subject B. The experimenter allotted 5 units of production to Subject A and 10 units to subject B. Then the transaction of allotted units of production was conducted by a certain rule (i.e., the Mitani mechanism). The allotted amounts corresponded to the reduction amounts in theory. In order to prepare an environment in which one subject (A) knew the production cost structure of the other subject (B), we explained the production cost to both subjects, and then conducted four practice rounds. Two were for subject A and two for subject B. Right before the real experiment, we announced who was subject A and who was B. Once the role of the subjects was determined, it remained fixed across 20 rounds.

Table 6 displays the total cost tables for subject A. The upper table is the payoff table for subject A. The payoff for subject A is determined by p_B, announced by subject B and the amount of transaction, z, by subject A without considering the charge term. If the prices announced by both subjects were different, subject A paid the charge. Subject A could also see the payoff table for subject B, which is shown in the bottom table in Table 6. The payoff for subject B is determined by p_A and z is announced by subject A. That is, subject B cannot change his or her own payoff by changing p_B. We will find the subgame perfect equilibrium through Table 6. Subject A first solves the optimization problem in stage 2 and then chooses a z that minimizes the total cost of subject A depending on the announcement of p_B by subject B. This is $z = z(p_B)$. For example, if $p_B = 6$, then $z = 1$. The diagonal from the upper left to the bottom right corresponds to $z = z(p_B)$ in Table 6. In stage 1, subject A should announce $p_A = 6$, since $p_B = 6$, to avoid the charge. However, these announcements are not a subgame perfect equilibrium. When $p_A = 6$, $z = 6$ makes the cost of subject

Table 6. The total cost table for Subject A under the Mitani Mechanism

	Your Choice of transaction (A)															
B's choice of price	–5	–4	–3	–2	–1	0	1	2	3	4	5	6	7	8	9	10
0	37.5	38	39.5	42	45.5	50	55.5	62	69.5	78	87.5	98	109.5	122	135.5	150
1	42.5	42	42.5	44	46.5	50	54.5	60	66.5	74	82.5	92	102.5	114	126.5	140
2	47.5	46	45.5	46	47.5	50	53.5	58	63.5	70	77.5	86	95.5	106	117.5	130
3	52.5	50	48.5	48	48.5	50	52.5	56	60.5	66	72.5	80	88.5	98	108.5	120
4	57.5	54	51.5	50	49.5	50	51.5	54	57.5	62	67.5	74	81.5	90	99.5	110
5	62.5	58	54.5	52	50.5	50	50.5	52	54.5	58	62.5	68	74.5	82	90.5	100
6	67.5	62	57.5	54	51.5	50	49.5	50	51.5	54	57.5	62	67.5	74	81.5	90
7	72.5	66	60.5	56	52.5	50	48.5	48	48.5	50	52.5	56	60.5	66	72.5	80
8	77.5	70	63.5	58	53.5	50	47.5	46	45.5	46	47.5	50	53.5	58	63.5	70
9	82.5	74	66.5	60	54.5	50	46.5	44	42.5	42	42.5	44	46.5	50	54.5	60
10	87.5	78	69.5	62	55.5	50	45.5	42	39.5	38	37.5	38	39.5	42	45.5	50
11	92.5	82	72.5	64	56.5	50	44.5	40	36.5	34	32.5	32	32.5	34	36.5	40
12	97.5	86	75.5	66	57.5	50	43.5	38	33.5	30	27.5	26	25.5	26	27.5	30
13	102.5	90	78.5	68	58.5	50	42.5	36	30.5	26	22.5	20	18.5	18	18.5	20
14	107.5	94	81.5	70	59.5	50	41.5	34	27.5	22	17.5	14	11.5	10	9.5	10
15	112.5	98	84.5	72	60.5	50	40.5	32	24.5	18	12.5	8	4.5	2	0.5	0

Table 6. (cont'd)

Your (A) choice of price	*Your Choice of transaction (A)*															
	−5	−4	−3	−2	−1	0	1	2	3	4	5	6	7	8	9	10
0	168.8	147	126.8	108	90.8	75	60.8	48	36.8	27	18.8	12	6.8	3	0.8	0
1	163.8	143	123.8	106	89.8	75	61.8	50	39.8	31	23.8	18	13.8	11	9.8	10
2	158.8	139	120.8	104	88.8	75	62.8	52	42.8	35	28.8	24	20.8	19	18.8	20
3	153.8	135	117.8	102	87.8	75	63.8	54	45.8	39	33.8	30	27.8	26	27.8	30
4	148.8	131	114.8	100	86.8	75	64.8	56	48.8	43	38.8	36	34.8	34	36.8	40
5	143.8	127	111.8	98	85.8	75	65.8	58	51.8	47	43.8	42	41.8	42	45.8	50
6	138.8	123	108.8	96	84.8	75	66.8	60	54.8	51	48.8	48	48.8	50	54.8	60
7	133.8	119	105.8	94	83.8	75	67.8	62	57.8	55	53.8	54	55.8	58	63.8	70
8	128.8	115	102.8	92	82.8	75	68.8	64	60.8	59	58.8	60	62.8	66	72.8	80
9	123.8	111	99.8	90	81.8	75	69.8	66	63.8	63	63.8	66	69.8	74	81.8	90
10	118.8	107	96.8	88	80.8	75	70.8	68	66.8	67	68.8	72	76.8	82	90.8	100
11	113.8	103	93.8	86	79.8	75	71.8	70	69.8	71	73.8	78	83.8	90	99.8	110
12	108.8	99	90.8	84	78.8	75	72.8	72	72.8	75	78.8	84	90.8	98	108.8	120
13	103.8	95	87.8	82	77.8	75	73.8	74	75.8	79	83.8	90	97.8	107	117.8	130
14	98.8	91	84.8	80	76.8	75	74.8	76	78.8	83	88.8	96	104.8	115	126.8	140
15	93.8	87	81.8	78	75.8	75	75.8	78	81.8	87	93.8	102	111.8	123	135.8	150

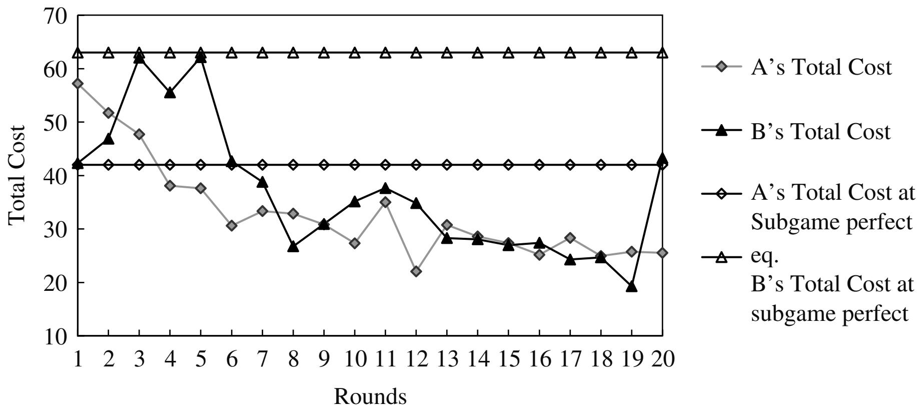

Figure 11. Average Total Costs when the charge is 10

B the minimum. Then, subject B would choose $p_B = 11$ since subject B incorporates the behavior of subject A, that is, $z = z(p_B)$. Therefore, $z = 1$ and $p_A = 6$ are not the best responses to subject A since subject A could avoid the charge by announcing $p_B = 11$. That is, $z = 1$ and $p_A = p_B = 6$ are not a subgame perfect equilibrium. Consider now that subject B announces $p_B = 9$. Then, subject A would choose $z = 4$ so that subject A could minimize his or her total cost. On the other hand, subject B would announce $p_A = 9$, which is the same as the announcement of subject A. Then, subject B would notice that $z = 4$ would minimize the total cost under $p_A = 9$. In order for subject A to choose $z = 4$, subject B announces $p_B = 9$ taking into account $z = z(p_B)$. That is, $z = 4$ and $p_A = p_B = 9$ are the subgame perfect equilibrium. The total cost is 42 for subject A and 63 for subject B.

Figure 11 shows the average total costs of subjects A and B when the charge is 10. They are smaller than those at subgame perfect equilibrium and they decrease with experience. We therefore make the following observation:

Observation 9. *When the charge is 10, the average total costs of subjects A and B are smaller than those at the subgame perfect equilibrium, they decreases with experience, and no pair who played subgame perfect equilibrium strategies was found.*

Why does the subject not adhering to subgame perfect equilibrium play? In early rounds, subjects noticed from Table 6 that a strategy profile of $z = 10$, $p_A = 0$, and $p_B = 15$ made subject A's cost 10 and subject B's cost 0. Under this strategy profile, subject A's cost is 0, but he or she must pay the charge since the two prices are not the same. Notice further that this profile is not a Nash equilibrium because subject A could avoid the charge by announcing $p_A = 15$. Consider the implication of this strategy profile. Subject A can make the purchasing price of emissions permits for subject B free of charge, and subject B can make the selling price of them for subject A as high as possible. Our highest price in Table 1 is 15. At the same time, the profile maximizes the number of transactions. Although subject A must pay the charge, the payoff profile of this strategy profile is strictly Pareto superior to the payoff profile at the subgame perfect equilibrium. We found 6 pairs who followed this strategy. Since the pair was not changed during 20 rounds, cooperation emerged. On the other hand, there were 2 pairs who converged to a Nash equilibrium. One pair's equilibrium was $z = 3$ and $p_A = p_B = 8$, and the other was $z = 2$ and $p_A = p_B = 7$. No pair played the subgame perfect equilibrium strategy.

When the charge was 50, 2 pairs reached an outcome where subject A's total cost was 50, and subject B's total cost was 0, which is different from the case when the charge was 10. Subject A in one of the two pairs chose $p_B = 15$ to make his or her charge zero. That is, subject A betrayed subject B. This seems an apparent effect of raising the charge. One pair converged to $z = 3$ and $p_A = p_B = 8$. No pair played the subgame perfect equilibrium strategy.

Summarizing the above, we make the following observation:

Observation 10. *When the charge is 50, the average total costs of subjects A and B are more than those at subgame perfect equilibrium, and no pair was found who played subgame perfect equilibrium strategies.*

In comparing these two sessions, consider first the choice of prices in stage 1. Two types of subject A's behavior were observed. One is cooperative behavior such that subject A chose a price as low as possible. If this is the case, subject A must bear the charge. In the second type, subject A chose the same price as subject B. In the charge 10 session, the former was mainly observed, and in the charge 50 session, the latter was mainly observed. As for the behavior of subject B, in the charge 10 session, subject B cooperated with subject A. In the charge 50 session, subject B tried to cooperate with subject A and make the total cost zero. But, most of the A subjects did not pay 50. The price distributions of the charge 10 and 50 sessions are shown in Figure 12.

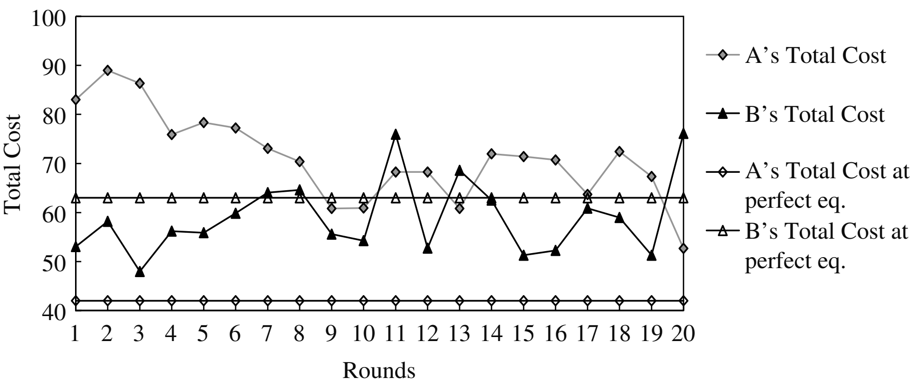

Figure 12. Average Total Costs when the charge is 50

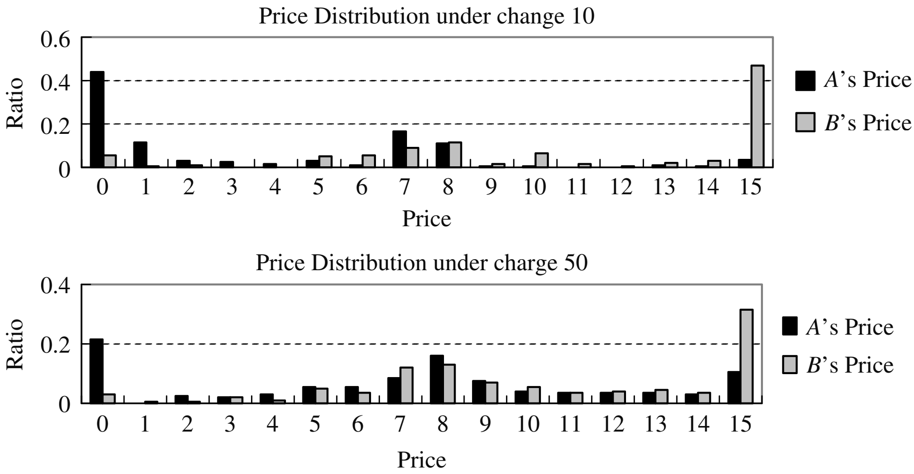

Figure 13. The Price Distribution of Charges 10 and 50

Subject *A* chose 0 and *B* chose 15 overwhelmingly in the case of charge 10 and the ratios go down in the case of charge 50. However, the ratios around 7 and 8 go up with charge 50.

Whether subjects understood the game or not is an important question. The ratio of the best response of subject *A* in stage 2 is 82%. That is, at least subject *A* seemed to understand the stage game.

The Mitani mechanism is a special case of the compensation mechanism by Varian (1994). The following observations are also applicable to this compensation mechanism. First, there are many Nash equilibria even though the subgame perfect equilibrium is unique. That is, subjects could not distinguish them. Second, subject *B*'s payoff would not be changed once subject *A*'s strategy is given. That is, whatever strategy subject *B* chooses, this does not affect his or her own payoff. The same problem was also found in the pivotal mechanism in the provision of public goods. This property might be the reason why the Mitani mechanism did not perform well in experiments. The third problem is the penalty scheme. Theoretically, the penalty should be zero when $p_A = p_B$ and positive when $p_A \neq p_B$. However, the special penalty scheme that the subjects employed might have influenced the results. It seems that the charge of 50 works slightly better than the charge of 10. That is, the shape of the penalty functions seems to be an important factor.

7. CONCLUDING REMARKS

The choice of a model is an important step in understanding how a specific economic phenomenon such as global warming works. We have reviewed three theoretical approaches, namely a simple microeconomic model, a social choice concept (i.e., strategy-proofness), and mechanisms constructed by theorists. The implicit environments on which the theories are based are quite different from one another, and theorists in varying fields may not realize the differences. Due to these differences, theories may result in contradictory conclusions. The social choice approach presents quite a negative view of attaining efficiency, but the two other approaches suggest some ways to attain it. From the point of view of policy makers, the environments conceived by theorists differ from the *real* environment that the policy makers must face. Unfortunately, we do not have any scientific measure of the differences between the environment of a theoretical model and the one of the real world.

A simple way to understand how each model works is to conduct experiments that implement the models' assumptions. The starting point is to create the environment in the experimental lab. If it works well, then the theory passes the experimental test. If not, the theory might have some flaw in its formulation. The failure of the test makes the policy makers look away. On the other hand, passing the experimental test does not necessarily mean that the policy maker should employ it.

For example, the experimental success of a model that does not include an explicit abatement investment decision should be compared with the experimental failure of a model with an explicit decision. The policy makers must consider the difference the environments that the theories are based upon.

The experimental approach helps us to draw conclusions on how and where theories work, and this approach is important for finding a real policy tool that can be used.

ACKNOWLEDGMENT

This study was partially supported by the Abe Fellowship, the Grant in Aid for Scientific Research 1143002 of the Ministry of Education, Science and Culture in Japan, the Asahi Glass Foundation, and the Nomura Foundation.

NOTES

[1] See Xepapadeas (1997) for standard theories on emissions trading.
[2] See also Schmalensee et al. (1998), Stavins (1998), and Joskow et al. (1998).
[3] See Kaino, Saijo, and Yamato (1999).
[4] The Mitani mechanism is based on a compensation mechanism proposed by Varian (1994).
[5] Saijo and Yamato (1999) consider an equilibrium when participation is a strategic variable.
[6] The same problem exists under social choice approach.

REFERENCES

Bohm, Peter, (June 1997). A Joint Implementation as Emission Quota Trade: An Experiment Among Four Nordic Countries, Nord 1997:4 by Nordic Council of Ministers.

Bohm, Peter, (January 2000). "Experimental Evaluations of Policy Instruments," mimeo.

Cason, Timothy N., (September 1995). "An Experimental Investigation of the Seller's Incentives in the EPA's Emission Trading Auction," *American Economic Review*, 85(4), pp. 905–22.

Cason, Timothy N. and Charles R. Plott, (March 1996). "EPA's New Emissions Trading Mechanism: A Laboratory Evaluation," *Journal of Environmental Economics and Management*, 30(2), pp. 133–60.

Dasgupta, Partha S., Peter J. Hammond, Eric S. Maskin, (April 1979). "The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility," *Review of Economic Studies*, 46(2), pp. 185–216.

Godby, Robert W., Stuart Mestelman, and R. Andrew Muller, (1998). "Experimental Tests of Market Power in Emission Trading Markets," in *Environmental Regulation and Market Structure*, Emmanuel Petrakis, Eftichios Sartzetakis, and Anastasios Xepapadeas (Eds.), Cheltenham, United Kingdom: Edward Elgar Publishing Limited.

Hizen, Yoichi, and Tatsuyoshi Saijo, (September 2001). "Designing GHG Emissions Trading Institutions in the Kyoto Protocol: An Experimental Approach," *Environmental Modelling and Software*, 16(6), pp. 533–543.

Hizen, Yoichi, Takao Kusakawa, Hidenori Niizwa and Tatsuyoshi Saijo, (November 2000). "GHG Emissions Trading Experiments: Trading Methods, Non-Compliance Penalty and Abatement Irreversibility."

Hurwicz, Leonid, (1979). "Outcome Functions Yielding Walrasian and Lindahl Allocations at Nash Equilibrium Points," *Review of Economic Studies*, 46, pp. 217–225.

Johansen, Leif, (Feb. 1977). "The Theory of Public Goods: Misplaced Emphasis?" *Journal of Public Economics*, 7(1), pp. 147–52.

Joskow, Paul L., Richard Schmalensee, and Elizabeth M. Bailey, (September 1998). "The Market for Sulfur Dioxide Emissions," *American Economic Review*, 88(4), pp. 669–685.

Kaino, Kazunari, Tatsuyoshi Saijo and Takehiko Yamato, (November 1999). "Who Would Get Gains from EU's Quantity Restraint on Emissions Trading in the Kyoto Protocol?"

Mitani, Satoshi, (January 1998). Emissions Trading: Theory and Experiment, Master's Thesis presented to Osaka University, (in Japanese).

Mitani, Satoshi, Tatsuyoshi Saijo, and Yasuyo Hamaguchi, (May 1998). "Emissions Trading Experiments: Does the Varian Mechanism Work?" (in Japanese).

Muller, R. Andrew and Stuart Mestelman, (June–August 1998). "What Have We Learned From Emissions Trading Experiments?" *Managerial and Decision Economics*, 19(4–5), pp. 225–238.

Saijo, Tatsuyoshi and Takehiko Yamato, (1999). "A Voluntary Participation Game with a Non-Excludable Public Good," *Journal of Economic Theory*, 84, pp. 227–242.

Stavins, Robert N., (Summer 1998). "What Can We Learn from the Grand Policy Experiment? Lessons from SO^2 Allowance Trading," *Journal of Economic Perspectives*, 12(3), pp. 69–88.

Schmalensee, Richard, Paul L. Joskow, A. Denny Ellerman, Juan Pablo Montero, and Elizabeth M. Bailey, (Summer 1998). "An Interim Evaluation of Sulfur Dioxide Emissions Trading," *Journal of Economic Perspectives*, 12(3), pp. 53–68.

Tietenberg, Tom, (1999). Environmental and Natural Resource Economics, Addison Wesley Longman.

Varian, H.R. (1994). "A Solution to the Problem of Externalities When Agents Are Well-Informed," *American Economic Review*, 84, pp. 1278–1293.

Xepapadeas, (1997). Anastasios, *Advanced Principles in Environmental Policy*, Edward Elgar.

Chapter 4

INTERNET CONGESTION: A LABORATORY EXPERIMENT

Daniel Friedman
University of California, Santa Cruz

Bernardo Huberman
Hewlett-Packard Laboratories

Abstract

Human players and automated players (bots) interact in real time in a congested network. A player's revenue is proportional to the number of successful "downloads" and his cost is proportional to his total waiting time. Congestion arises because waiting time is an increasing random function of the number of uncompleted download attempts by all players. Surprisingly, some human players earn considerably higher profits than bots. Bots are better able to exploit periods of excess capacity, but they create endogenous trends in congestion that human players are better able to exploit. Nash equilibrium does a good job of predicting the impact of network capacity and noise amplitude. Overall efficiency is quite low, however, and players overdissipate potential rents, i.e., earn lower profits than in Nash equilibrium.

1. INTRODUCTION

The Internet suffers from bursts of congestion that disrupt cyberspace markets. Some episodes, such as gridlock at the Victoria's Secret site after a Superbowl advertisement, are easy to understand, but other episodes seem to come out of the blue. Of course, congestion is also important in many other contexts. For example, congestion sometimes greatly degrades the value of freeways, and in extreme cases (such as burning nightclubs) congestion can be fatal. Yet the dynamics of congestion are still poorly understood, especially when (as on the Internet) humans interact with automated agents in real time.

In this paper we study congestion dynamics in the laboratory using a multiplayer interactive video game called StarCatcher. Choices are real-time (i.e., asynchronous):

A. Rapoport and R. Zwick (eds.), Experimental Business Research, Vol. II, 83–102.

at every instant during a two minute period, each player can start to download or abort an uncompleted download. Human players can freely switch back and forth between manual play and a fully automated strategy. Other players, called bots, are always automated. Players earn revenue each time they complete the download, but they also accumulate costs proportional to waiting time.

Congestion arises because waiting time increases stochastically in the number of pending downloads. The waiting time algorithm is borrowed from Maurer and Huberman (2001), who simulate bot-only interactions. This study and earlier studies show that congestion bursts arise from the interaction of many bots, each of whom reacts to observed congestion observed with a short lag. The intuition is that bot reactions are highly correlated, leading to non-linear bursts of congestion.

At least two other strands of empirical literature relate to our work. Ochs (1990), Rapoport et al. (1998) and others find that human subjects are remarkably good at coordinating entry into periodic (synchronous) laboratory markets subject to congestion. More recently, Rapoport et al. (2003) and Seale et al. (2003) report fairly efficient queuing behavior in a laboratory game that has some broad similarities to ours, but (as discussed in section 5 below) differs in numerous details.

A separate strand of literature considers asynchronous environments, sometimes including bots. The Economist (2002) mentions research by Dave Cliff at HP Labs Bristol intended to develop bots that can make profits in major financial markets that allow asynchronous trading. The article also mentions the widespread belief that automated trading strategies provoked the October 1987 stock market crash. Eric Friedman et al. (forthcoming) adapt periodic laboratory software to create a near-asynchronous environment where some subjects can update choices every second; other subjects are allowed to update every 2 seconds or every 30 seconds. The subjects play quantity choice games (e.g., Cournot oligopoly) in a very low information environment: they know nothing about the structure of the payoff function or the existence of other players. Play tends to converge to the Stackelberg equilibrium (with the slow updaters as leaders) rather than to the Cournot equilibrium. In our setting, by contrast, there is no clear distinction between Stackelberg and Cournot, subjects have asynchronous binary choices at endogenously determined times, and they compete with bots.

After describing the laboratory set up in the next section, we sketch theoretical predictions derived mainly from Nash equilibrium. Section 4 presents the results of our experiment. Surprisingly, some human players earn considerably higher profits than bots. Bots are better able to exploit periods of excess capacity, but they create endogenous trends in congestion that human players are better able to exploit. The comparative statics of pure strategy Nash equilibrium do a good job of predicting the impact of network capacity and noise amplitude. However, overall efficiency is quite low relative to pure strategy Nash equilibrium, i.e., players "overdissipate" potential rents.

Section 5 offers some perspectives and suggestions for follow up work. Appendix A collects the details of algorithms and mathematical derivations. Appendix B reproduces the written instructions to human subjects.

2. THE EXPERIMENT

The experiment was conducted at UCSC's LEEPS lab. Each session lasts about 90 minutes and employs at least four human subjects, most of them UCSC undergraduates. Students sign up on line after hearing announcements in large classes, and are notified by email about the session time and place, using software developed by UCLA's CASSEL lab. Subjects read the instructions attached in Appendix B, view a projection of the user interface, participate in practice periods, and get public answers to their questions. Then they play 16 or more periods of the StarCatcher game. At the end of the session, subjects receive cash payment, typically \$15 to \$25. The payment is the total points earned in all periods times a posted payrate, plus a \$5.00 show-up allowance.

Each StarCatcher period lasts 240 seconds. At each instant, any idle player can initiate a service request by clicking the Download button, as in Figure 1. The service delay, or latency λ, is determined by an algorithm sketched in the paragraph after next. Unless the download is stopped earlier, after λ seconds the player's screen flashes a gold star and awards her 10 points. However, each second of delay costs the player 2 points, so she loses money on download requests with latencies greater than 5 seconds. The player can't begin a second download while an earlier request is still being processed but she can click the Stop button; to prevent excessive losses the computer automatically stops a request after 10 seconds. The player can also click the Reload button, which is equivalent to Stop together with an immediate new download request, and can toggle between manual mode (as just described) and automatic mode (described below).

The player's timing decision is aided by a real-time display showing the results of all service requests terminating in the previous 10 seconds. The player sees the mean latency as well as a latency histogram that includes Stop orders, as illustrated in Figure 1.

The delay algorithm is a noisy version of a single server queue model known in the literature as M/M/1. Basically, the latency λ is proportional to the reciprocal of current idle capacity. For example, if capacity is 6 and there are currently 4 active users, then the delay is proportional to $1/(6-4) = 1/2$. In this example, 5 users would double the delay and 6 users would make the delay arbitrarily long. As explained in Appendix A, the actual latency experienced by a user is modified by a mean reverting noise factor, and is kept positive and finite by truncating at specific lower and upper bounds.

The experiments include automated players (called robots or bots) as well as humans. The basic algorithm for bots is: initiate a download whenever the mean latency (shown on all players' screens) is less than 5 seconds minus a tolerance, i.e., whenever it seems sufficiently profitable. The tolerance averages 0.5 seconds, corresponding to an intended minimum profit margin of 1 point per download. Appendix A presents details of the algorithm. Human players in most sessions have the option of "going on autopilot" using this algorithm, as indicated by the toggle button in Figure 1 Go To Automatic / Go To Manual. Subjects are told,

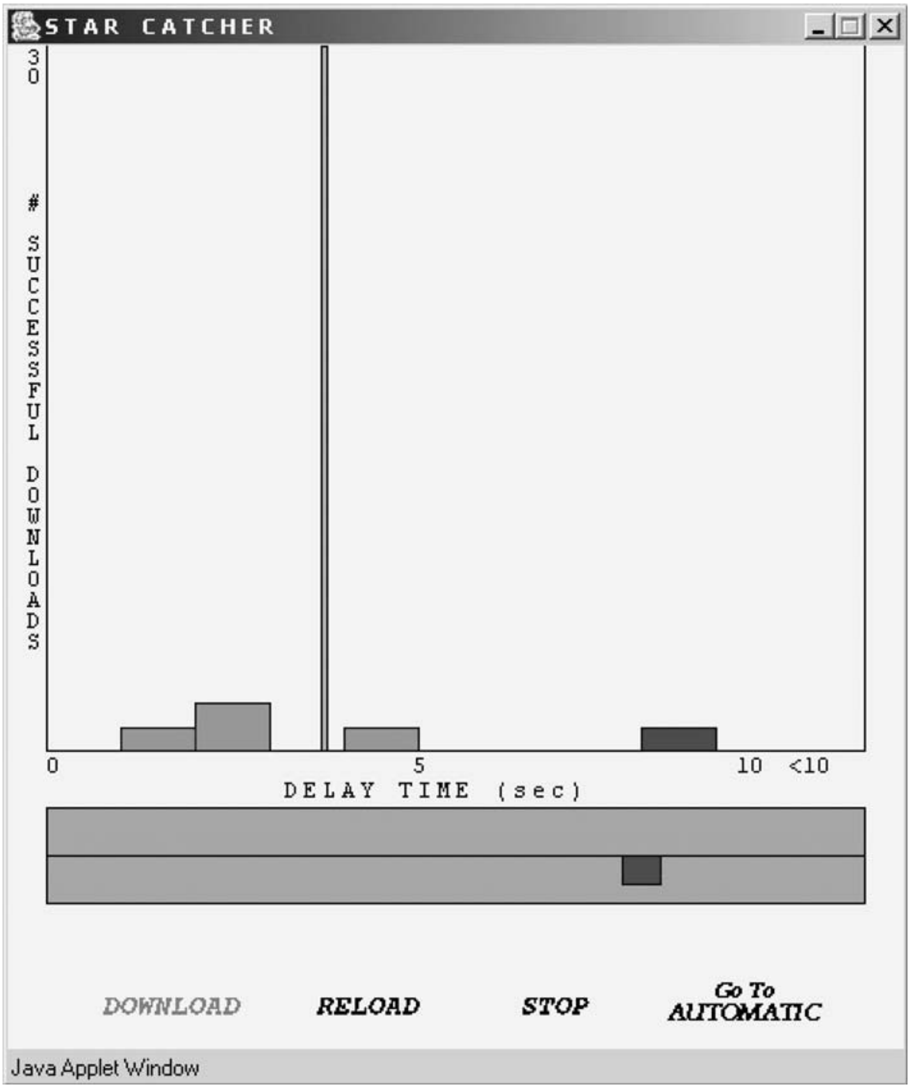

Figure 1. User interface. The four decision buttons appear at the bottom of the screen; the download button is faded because the player is currently waiting for his download request to finish. The dark box on the thick horizontal bar just above the decision buttons indicates an 8 second waiting time (hence a net loss) so far. The histogram above reports results of download requests from all players terminating in the last 10 seconds. Here, one download took 2 seconds, two took 3 seconds, one took 5 seconds and one took 9 seconds. The color of the histogram bar indicates whether the net payoff from the download was positive (green, here light grey) or negative (red, here dark grey). The thin vertical line indicates the mean delay, here about 3.7 seconds. Time remaining is shown in a separate window.

> When you click GO TO AUTOMATIC a computer algorithm decides for you when to download. There sometimes are computer players (in addition to your fellow humans) who are always in AUTOMATIC. The algorithm mainly looks at the level of recent congestion and downloads when it is not too large.

The network capacity and the persistence and amplitude of the background noise is controlled at different levels in different periods. The number of human players and bots also varies; the humans who are sidelined from StarCatcher for a few periods use the time to play an individual choice game such as TreasureHunt, described in Friedman et al. (2003). Table 1 summarizes the values of the control variables used in all sessions analyzed below.

3. THEORETICAL PREDICTIONS

A player's objective each period is to maximize profit $\Pi = rN - cL$, where r is the reward per successful download, N is the number of successful downloads, c is the delay cost per second, and L is the total latency time summed over all download attempts in that period. The relevant constraints include the total time T in the period, and the network capacity C. The constant of proportionality for latency, i.e., the time scale S, is never varied in our experiments.

An important benchmark is social value V^*, the maximized sum of players' profits. That is, V^* is the maximum total profit obtainable by an omniscient planner who controls players' actions. Appendix A shows that, ignoring random noise, that benchmark is given by the expression $V^* = 0.25S^{-1}Tr(1 + C - cS/r)^2$. Typical parameter values in the experiment are $T = 120$ seconds, $C = 6$ users, $S = 8$ user-sec, $c = 2$ points/sec and $r = 10$ points. The corresponding social optimum values are $U^* = 2.70$ active users, $\lambda^* = 1.86$ seconds average latency, $\pi^* = 6.28$ points per download, $N^* = 174.2$ downloads, and $V^* = 1094$ points per period.

Of course, a typical player tries to increase his own profit, not social value. A selfish and myopic player will attempt to download whenever the incremental apparent profit π is sufficiently positive, i.e., whenever the reward $r = 10$ points sufficiently exceeds the cost λc at the currently displayed average latency λ. Thus such a player will choose a latency threshold ε and follow

Rule R. If idle, initiate a download whenever $\lambda \leq r/c - \varepsilon$.

In Nash equilibrium (NE) the result typically will be inefficient congestion, because an individual player will not recognize the social cost (longer latency times for everyone else) when choosing to initiate a download. Our game has many pure strategy NE due to the numerous player permutations that yield the same overall outcome, and due to integer constraints on the number of downloads. Fortunately, the NE are clustered and produce outcomes in a limited range.

To compute the range of total NE total profit V^{NE} for our experiment, assume that all players use the threshold $\varepsilon = 0$ and assume again that noise is negligible. No

Table 1. Design of Sessions

Date	*# of periods*	*# of player-periods*										*Max # of robots*	*Max # human players*	*Experienced humans*
		Total	*By volatility*		*By capacity*									
			low	*high*	*2*	*3*	*4*	*5*	*6*	*7*	*9*			
8/21/02	27	159	87	72	20	101	38	0	0	0	0	4	4	no
8/22/02	32	189	94	95	19	120	50	0	0	0	0	4	4	no
8/20/02	32	192	97	95	21	117	54	0	0	0	0	4	4	yes
9/11/02	32	243	130	113	0	56	90	0	97	0	0	0	6	yes
9/12/02	32	199	101	98	20	126	53	0	0	0	0	5	3	yes
9/5/02	32	193	99	94	20	121	52	0	0	0	0	4	4	no
1/24/03	16	127	54	73	0	37	40	0	50	0	0	4	6	no
1/31/03	24	155	77	78	20	105	30	0	0	0	0	4	4	no
2/5/03	27	216	120	96	0	54	72	0	90	0	0	4	6	yes
2/4/03	16	104	52	52	10	64	30	0	0	0	0	4	4	no
2/12/03	18	143	71	72	0	36	47	0	60	0	0	4	6	no
2/14/03	27	214	119	95	0	54	72	0	88	0	0	4	6	no
2/19/03	31	194	100	94	19	123	52	0	0	0	0	4	4	yes
5/23/03	27	164	89	75	0	94	64	6	0	0	0	5	6	no
10/2/03	31	112	63	49	76	22	6	8	0	0	0	4	4	no
10/3/03	27	164	86	78	17	45	50	8	0	20	24	6	6	no

Volatility: low: Sigma = .0015, Tau = .0002; Volatility: high: Sigma = .0025, Tau = .00002 player will earn negative profits in NE, since the option is always available to remain idle and earn zero profit. Hence the lower bound on V^{NE} is zero. Appendix A derives the upper bound $V^{MNE} = T(rC - cS)/S$ from the observation that it should never be possible for another player to enter and earn positive profits. Hence the maximum NE efficiency is $V^{MNE}/V^* = 4(C - cS/r)/(1 + C - cS/r)^2 = 4U^{MNE}/(1 + U^{MNE})^2$. For the parameter values used above ($T = 120$, $C = 6$, $S = 8$, $c = 2$ and $r = 10$), the upper bound NE values are $U^{MNE} = 4.4$ active users (players), $\lambda^{MNE} = 3.08$ seconds delay, $\pi^{MNE} = 3.85$ points per download, $N^{MNE} = 171.6$ downloads, and $V^{MNE} = 660.1$ points per period, for a maximum efficiency of 60.4%.

The preceding calculations assume that the number of players m in the game is at least $U^{MNE} + 1$, so that congestion can drive profit to zero. If there are fewer players, then in Nash equilibrium everyone is always downloading. In this case there is excess capacity $a = U^{MNE} + 1 - m = C + 1 - cS/r - m > 0$ and, as shown in the Appendix, the interval of NE total profit shrinks to a single point, $\Pi m = Tram/S$.

What happens if the background noise is not negligible? As explained in the Appendix, the noise is mean-reverting in continuous time. Thus there will be some good times when effective capacity is above C and some bad times when it is lower. Since the functions V^{MNE} and V^* are convex in C (and bounded below by zero), Jensen's inequality tells us that the loss of profit in bad times does not fully offset the gain in good times. When C and m are sufficiently large (namely, $m > C > cS/r + 1$, where the last expression is 2.6 for the parameters above), this effect is stronger for V^* than for V^{MNE}. In this case Nash equilibrium efficiency V^{MNE}/V^* decreases when there is more noise. Thus the prediction is that aggregate profit should increase but that efficiency should decrease in the noise amplitude $\sigma/\sqrt{2\tau}$ (see Appendix A).[1]

A key testable prediction arises directly from the Nash equilibrium benchmarks. The null hypothesis, call it full rent dissipation, is that players' total profits will be in the Nash equilibrium range. That is, when noise amplitude is small, aggregate profits will be $V^{MNE} = Tram/S$ in periods with excess capacity $a > 0$, and will be between 0 and $V^{MNE} = T(rC - cS)/S$ in periods with no excess capacity. The corresponding expressions for efficiency have already been noted.

One can find theoretical support for alternative hypotheses on both sides of the null. Underdissipation refers to aggregate profits higher than in any Nash equilibrium, i.e., above V^{MNE}. This would arise if players can maintain positive thresholds ε in Rule R, for example. A libertarian justification for the underdissipation hypothesis is that players somehow self-organize to partially internalize the congestion externality (see e.g., Gardner, Ostrom, and Walker, 1992). For example, players may discipline each other using punishment strategies. Presumably the higher profits would emerge in later periods as self-organization matures. An alternative justification from behavioral economics is that players have positive regard for the other players' utility of payoffs, and will restrain themselves from going after the last penny of personal profits in order to reduce congestion. One might expect this effect to weaken a bit in later periods.

Overdissipation of rent, i.e., negative aggregate profits, is the other possibility. One theoretical justification is that players respond to relative payoff and see increasing returns to downloading activity (e.g., Hehenkamp et al., 2001). A behavioral economics justification is that people become angry at the greed of other players and are willing to pay the personal cost of punishing them by deliberately increasing congestion (e.g., Cox and Friedman, 2002). Behavioral noise is a third possible justification. For example, Anderson, Goeree and Holt (1998) use quantal response equilibrium, in essence Nash equilibrium with behavioral noise, to explain overdissipation in all-pay auctions.

Further insights may be gained from examining individual decisions. The natural null hypothesis is that human players follow Rule R with idiosyncratic values of the threshold ε. According to this hypothesis, the only significant explanatory variable for the download decision will be $\lambda - r/c = \lambda - 5$ sec, where λ is the average latency currently displayed on the screen. An alternative hypothesis (which occurred to us only after looking at the data) is that some humans best-respond to Rule R behavior, by anticipating when such behavior will increase or decrease λ and reacting to the anticipation.

The experiment originally was motivated by questions concerning the efficiency impact of automated Rule R strategies. The presumption is that bots (and human players in auto mode) will earn higher profits than humans in manual mode.[2] How strong is this effect? On the other hand, does a greater prevalence of bots depress everyone's profit? If so, is the second effect stronger than the first, i.e., are individual profits lower when everyone is in auto mode than when everyone is in manual mode? The simulations reported in Maurer and Huberman (2001) confirm the second effect but disconfirm the social dilemma embodied in the last question. Our experiment examines whether human subjects produce similar results.

4. RESULTS

We begin with a qualitative overview of the data. Figure 2 below shows behavior in a fairly typical period. It is not hard to confirm that bots indeed follow the variable λ = average delay: their download requests cease when λ rises above 4 or 5, and the line indicating the number of bots downloading stops rising. It begins to decline as existing downloads are completed. Likewise, when λ falls below 4 or 5, the number of bot downloads starts to rise.

The striking feature about Figure 2 is that the humans are different. They appear to respond as much to the *change in* average delay. Sharp decreases in average delay encourage humans to download. Perhaps they anticipate further decreases, which would indeed be likely if most players use Rule R. We shall soon check this conjecture more systematically.

Figure 3 shows another surprise, strong overdissipation. Both bots and humans lose money overall, especially bots (which include humans in the auto mode). The top half of human players spend only 1% of their time in auto mode, and even the bottom half spend only 5% of their time in auto mode. In manual mode, bottom half human players lose lots of money but at only 1/3 the rate of bots, and top half humans actually make modestly positive profit.

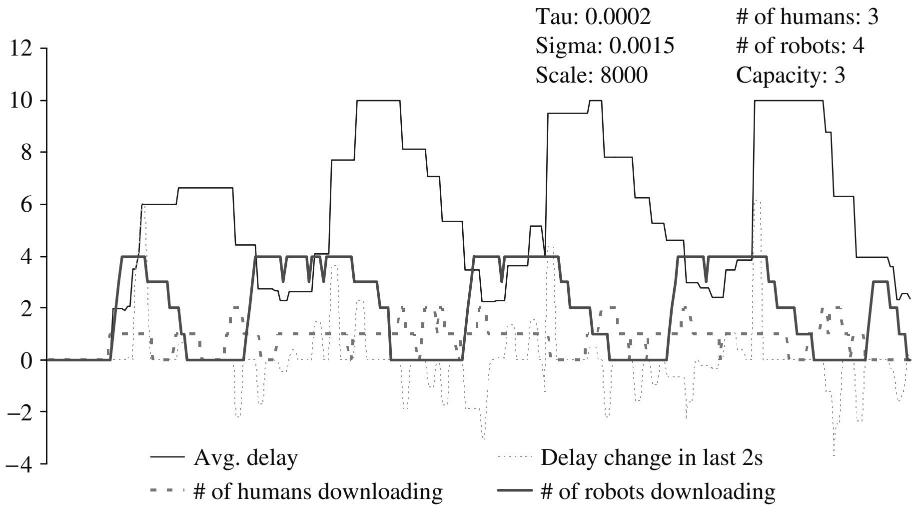

Figure 2. Exp. 09-12-2002, Period 1.

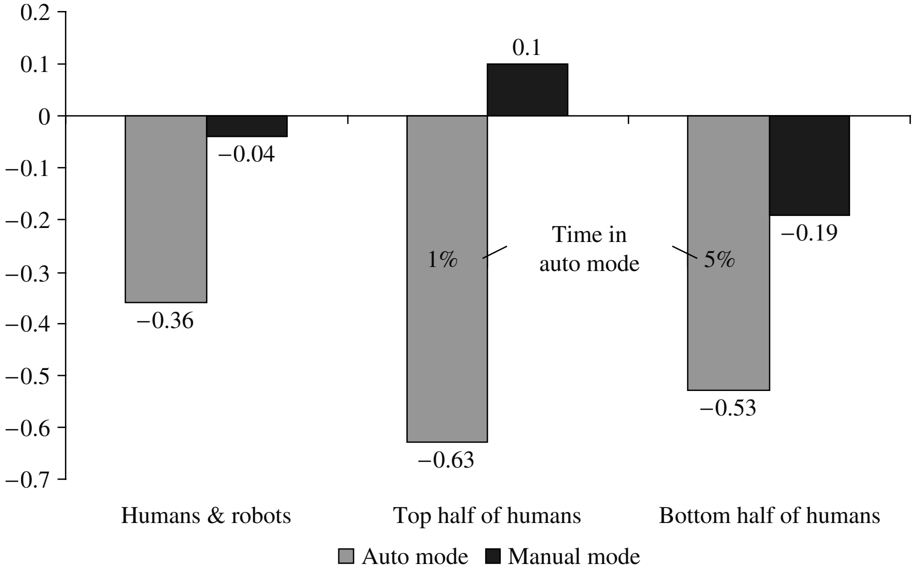

Figure 3. Profit per second in auto and manual mode.

Figure 4 offers a more detailed breakdown. When capacity is small, there is only a small gap between social optimum and the upper bound aggregate profit consistent with Nash Equilibrium, so Nash efficiency is high as shown in the green bars for $C = 2, 3, 4$. Bots lose money rapidly in this setting because congestion sets in

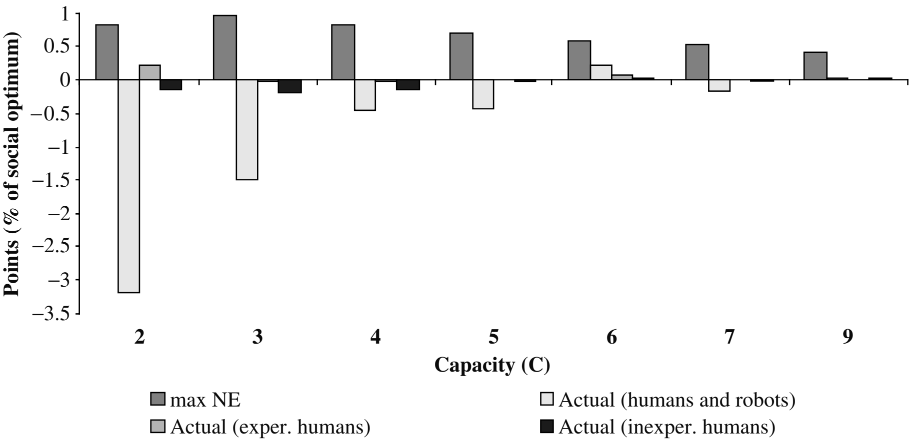

Figure 4. Theoretical and actual profits as percentage of social optimum.

quickly when capacity is small. Humans lose money when inexperienced. Experienced human players seem to avoid auto mode and learn to anticipate the congestion sufficiently to make positive profits. When capacity is higher ($C = 6$), bots do better even than experienced humans, perhaps because they are better at exploiting the good times with excess capacity. (Of course, overdissipation is not feasible with excess capacity: in NE everyone downloads as often as physically possible and everyone earns positive profit.)

We now turn to more systematic tests of hypotheses. Table 2 below reports OLS regression results for profit rates (net payoff per second) earned by four types of players. The first column shows that bots (lumped together with human players in auto mode) do much better with larger capacity and with higher noise amplitude, consistent with NE predictions. The effects are highly significant, statistically as well as economically. The other columns indicate that humans in manual mode are able to exploit increases in capacity only about half as much as bots, although the effect is still statistically highly significant for all humans and top half of humans. The next row suggests that bots but not humans are able to exploit higher amplitude noise. The last row of coefficient estimates finds that, in our mixed bot-human experiments, the interaction [noise amplitude with excess fraction of players in auto mode] has the opposite effect for bots as in Maurer and Huberman (2001), and has no significant effect for humans.

Table 3 above reports a fine-grained analysis of download decisions, the dependent variable in the logit regressions. Consistent with Rule R (hardwired into their algorithm), the bots respond strongly and negatively to the average delay observed on the screen minus $r/c = 5$. Surprisingly, the regression also indicates that bots are more likely to download when the observed delay increased over the last 2 seconds; we interpret this as an artifact of the cyclical congestion patterns. Fortunately

Table 2. OLS Estimates of profit rates

Indep. variables	*For:*	*Auto mode All players*	*All humans*	*Manual mode Top half*	*Bottom half*
Intercept		0.88	0.48	0.64	not sig.
Excess capacity		0.69	0.27	0.29	0.17[a]
Excess capacity^2		0.08	0.03	0.04	not sig.
Noise		0.53	not sig.	not sig.	–0.12[b]
Noise*(s – 1/2)		–1.81	not sig.	not sig.	not sig.
NOBS		1676	1222	640	582

Notes: all significant at $p < 0.01$, except a: $p = 0.04$, b: $p = 0.06$
Excess capacity = $a = C - m - 0.6$
Noise = sigma/$\sqrt{(2\text{Tau})}$
s = fraction of all players in auto mode per period

Table 3. Logit regression for download decision

Indep. Variables	*For:*	*Auto mode All players*		*All humans*	*Manual mode Top half*	*Bottom half*
Intercept		–2.08	–2.12	–2.51	–2.31	–2.74
Avg. delay –5s		–0.31	–0.31	0.03	0.06	–0.01[a]
2s change in avg. delay		0.08		–0.19	–0.30	–0.08
NOBS		163,602	163,602	186,075	92,663	93,412

Notes: all significant at $p < 0.01$, except a: $p = 0.047$

the delay coefficient estimate is unaffected by omitting the variable for change in delay.

Human players react in the opposite direction to delay changes. The regressions confirm the impression gleaned from Figure 2 that humans are much more inclined to initiate download requests when the observed delay is decreasing. Perhaps surprisingly, experienced humans are somewhat more inclined to download when the observed delay is large. A possible explanation is that they then anticipate less congestion from bots.

The results reported above seem fairly robust to changes in the specification. In particular, including time trends within or across periods seems to have little systematic impact.

5. DISCUSSION

The most surprising result is that human players outperform the current generation of automated players (bots). The bots do quite badly when capacity is low. Their decision rule fails to anticipate the impact of other bots and neglects the difference between observed congestion (for recently completed download attempts) and anticipated congestion (for the current download attempt). Human players are slower and less able to exploit excess capacity (including transient episodes due to random noise), but some humans are far better at anticipating and exploiting the congestion trends that the bots create. In our experiment the second effect outweighs the first, so humans earn higher profits overall than bots.

Perhaps the most important questions in our investigation concerned rent dissipation. Would human players find some way to reduce congestion costs and move towards the social optimum, or would they perhaps create even more congestion than in Nash equilibrium? Sadly, overdissipation outcomes are most prevalent in our data.

The Nash comparative statics, on the other hand, generally help explain the laboratory data. Nash equilibrium profit increases in capacity and noise amplitude, and so do observed profits.

Several directions for future research suggest themselves. First, one might want to look at smarter bots. Preliminary results show that it is not as easy as we thought to find more profitable algorithms; linear extrapolation from available data seems rather ineffective. That project contemplates higher levels of sophistication (in a sense similar to Stahl and Wilson, 1995) but the results are not yet in.

Second, one might want to connect our research to the experiments on queuing behavior. As noted in the introduction, Rapoport et al. (2003) and a companion paper reported fairly efficient outcomes, rather different than our own. Which design differences from ours are crucial? The list of possible suspects is quite long: no bots; synchronous decisions in discrete time; a single service request per player each period; simultaneous choice at the beginning of the period; precommited requests (no counterpart to our "stop" or "reload"); deterministic and constant service times in a first-in, first-out queue; no information feedback during the period; and no information feedback between periods regarding congestion at times not chosen. Answering the question may not be easy, but it surely would be interesting.

More generally, one might want to probe the robustness of the overdissipation result. It clearly should be checked in humans-only and in bots-only environments, and preliminary results seem consistent with the findings reported above. One should also check alternative congestion functions to the mean-reverting noisy M/M/1 queuing process. Finally, it would be quite interesting to investigate mechanisms such as congestion taxes to see whether they enable humans and robots to earn healthier profits in congestible real-time environments.

ACKNOWLEDGMENT

The National Science Foundation funded this work under grant IIS-9986651. We are also grateful for the research assistance of Alessandra Cassar, Garrett Milam, Ralf Hepp, and Kai Pommerenke, and the programming support of Neil Macneale, Jeremy Avnet and Nitai Farmer. We benefited from comments by colleagues including Joshua Aizenman, Eileen Brooks, Rajan Lukose, Nirvikar Singh, and Donald Wittman, and by HKUST conference participants, especially Yan Chen, Rachel Croson and Eric Johnson. Significant improvements to the final version are due to the comments of an anonymous referee.

NOTES

[1] The simulations reported in Maurer and Huberman (2001) suggest an alternative hypothesis: profits increase in noise amplitude times ($s - 1/2$), where s is the fraction of players in auto mode. It should be noted that their bot algorithm supplemented Rule R with a Reload option.

[2] Indeed, a referee of our grant proposal argued that it was redundant to use human subjects. He thought it obvious that the bots would perform better.

[3] This variable is generated by summing up the times for successful (the download took less than or exactly ten seconds) and unsuccessful (failed download attempt, i.e., no download within ten seconds) download attempts that were *completed within the last ten seconds*. The result is then divided by the number of download attempts to lead to the average delay (*AD*). The variable is continuously updated. Times for download attempts that have been aborted (by the player hitting the "STOP" or the "RELOAD" button) are disregarded.

REFERENCES

Anderson, S., Goeree, J. and Holt, C., (August 1998). "The All-Pay Auction: Equilibrium with Bounded Rationality." *Journal of Political Economy*, 106(4), 828–853.

Cox J. C. and Friedman, D. (October 2002). "A Tractable Model of Reciprocity and Fairness," UCSC Manuscript.

Feller, William, (1968). An Introduction to Probability Theory and Its Applications, Vol 2. NY: Wiley.

Friedman, Eric, Mikhael Shor, Scott Schenker, and Barry Sopher, (November 30, 2002). "An Experiment on Learning with Limited Information: Nonconvergence, Experimentation Cascades, and the Advantage of Being Slow." *Games and Economic Behavior* (forthcoming).

Economist magazine, "Robo-traders," Nov. 30, 2002, p. 65.

Gardner, Roy, Ostrom, Elinor and Walker, James, (June 1992). "Covenants With and Without a Sword: Self-Governance is Possible." *American Political Science Review*, 86(2), 404–417.

Hehenkamp, Burkhard, Leininger, Wolfgang, and Possajennikov, Alex, (December 2001). "Evolutionary Rent Seeking." CESifo Working Paper 620.

Maurer, Sebastian and Bernardo Huberman, (2001). "Restart Strategies and Internet Congestion." *Journal of Economic Dynamics & Control* 25, 641–654.

Ochs, Jack, (May, 1990). "The Coordination Problem in Decentralized Markets: An Experiment." *The Quarterly Journal of Economics*, 105(2), 545–559.

Rapoport, A., Seale, D. A., Erev, I., & Sundali, J. A., (1998). "Equilibrium Play in Large Market Entry Games." *Management Science*, 44, 119–141.

Rapoport, A., Stein, W., Parco, J. and Seale, D., (July 2003). "Equilibrium Play in Single Server Queues with Endogenously Determined Arrival Times." University of Arizona Manuscript.

Seale, D., Parco, J., Stein, W. and Rapoport, A., (January 2003). "Joining a Queue or Staying Out: Effects of Information Structure and Service Time on Large Group Coordination." University of Arizona Manuscript.

Stahl, D. O. and Wilson, P., (1995). "On Players' Models of Other Players – Theory and Experimental Evidence." *Games and Economic Behavior*, 10, 213–254.

APPENDIX A. TECHNICAL DETAILS.

A.1. Latency and Noise. Following the noisy M/M/1 queuing model of Maurer and Huberman (2001), latency for a download request initiated at time t is

$$\lambda(t) = \frac{S[1 + e(t)]_+}{1 + C - U(t)} \tag{A1}$$

if the denominator is positive, and otherwise is $\lambda^{max} > 0$. To unpack the expression (A1), note that the subscripted "+" refers to the positive part, i.e., $[x]_+ = \max\{x, 0\}$. The parameter C is the capacity chosen for that period; more precisely, to remain consistent with conventions in the literature, C represents full capacity minus 1. The parameter S is the time scale, or constant of proportionality, and $U(t)$ is usage, the number of downloads initiated but not yet completed at time t. The experiment truncates the latency computed from (A1) to the interval [0.2, 10.0] seconds. The lower truncation earns the 10 point reward but the upper truncation at $\lambda^{max} = 10$ seconds does not.

The random noise $e(t)$ is Normally distributed with volatility σ and unconditional mean 0. The noise is mean reverting in continuous time and follows the Ornstein-Uhlenbeck process with persistence parameter $\tau > 0$ (see Feller, p. 336). That is, $e(0) = 0$ and, given the previous value $x = e(t - h)$ drawn at time $t - h > 0$, the algorithm draws a unit Normal random variate z and sets $e(t) = x\exp(-\tau h) + z\sigma\sqrt{[1 - \exp(-2\tau h)]/(2\tau)}$. Thus the conditional mean of noise is typically different from zero; it is the most recently observed value x shrunk towards zero via an exponential term that depends on the time lag h since the observation was made and a shrink rate $\tau > 0$. In the no-persistence (i.e., no mean reversion or shrinking) limit $\tau \to 0$, we have Brownian motion with conditional variance $\sigma^2 h$, and $e(t) = x + z\sigma\sqrt{h}$. In the long run limit as $h \to \infty$ we recover the unconditional variance $\sigma^2/(2\tau)$. The appropriate measure of noise amplitude in our setting therefore is its square root $\sigma/\sqrt{2\tau}$.

In our experiments we used two levels each for σ and τ. Rescaling time in seconds instead of milliseconds, the levels are 2.5 and 1.5 for σ, and 0.2 and 0.02 for τ. Figure A1 shows typical realizations of the noise factor $[1 + e(t)]_+$ for the two combinations used most frequently, low amplitude (low σ, high τ) and high amplitude (high σ, low τ).

A.2. Efficiency, no noise case. Social value V is the average net benefit $\pi = r - \lambda c$ per download times the total number of downloads $n \approx UT/\lambda$, where λ is the average latency, T is the length of a period and U is the average number of users attempting to download. Assume that $\sigma = 0$ (noise amplitude is zero) so by (A1) the average latency is $\lambda = S/(1 + C - U)$. Assume also that the expression for n is exact. Then the first order condition (taking the derivative of $V = \pi n$ with respect to U and finding the root) yields $U^* = 0.5(1 + C - cS/r)$. Thus $\lambda^* = 2S/(1 + C + cS/r)$, and so maximized social value is $V^* = 0.25S^{-1}Tr(1 + C - cS/r)^2$.

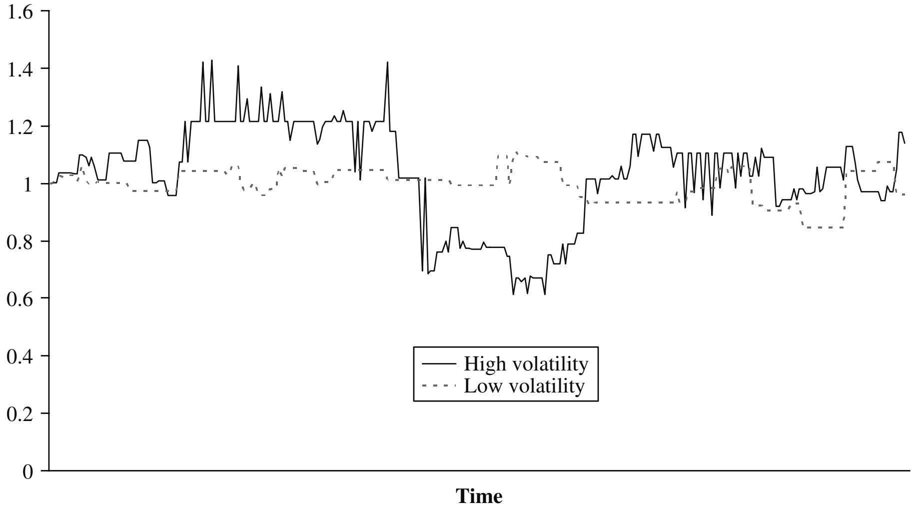

Figure A1. Noise (exp 10/3/03, periods 1 and 4).

To obtain the upper bound on social value consistent with Nash equilibrium, suppose that more than 10 seconds remain, the player currently is idle and the expected latency for the current download is λ. The zero-profit latency is derived from $0 = \pi = r - \lambda c$. Now $\lambda = r/c$ and the associated number of users is $U^{**} = 2U^* = C + 1 - cS/r$. Hence the minimum number of users consistent with NE is $U^{MNE} = U^{**} - 1 = C - cS/r$. The associated latency is $\lambda^{MNE} = rS/(r + cS)$, and the associated profit per download is $\pi^{MNE} = r^2/(r + cS)$, independent of C. The maximum number of downloads is $N^{MNE} = TU^{MNE}/\lambda^{MNE} = T(r + cS)(rC - cS)/(r^2S)$. Hence the upper bound on NE total profit is $V^{MNE} = N^{MNE}\pi^{MNE} = T(rC - cS)/S$, and the maximum NE efficiency is $V^{MNE}/V^* = (C - cS/r)/(1 + C - cS/r)^2 = 4U^{MNE}/(1 + U^{MNE})^2 \equiv Y$. Since $dU^{MNE}/dC = 1$, it follows that $dY/dC < 0$ iff $dY/dU^{MNE} < 0$ iff $1 < U^{MNE} = C - cS/r$. It is easy to verify that Y is $0(1/C)$.

A.3 Bot algorithm. In brief, the bot algorithm uses Rule R with a random threshold ε drawn independently from the uniform distribution on [0, 1.0] sec. The value of λ is the mean reported in the histogram window, i.e., the average for download requests completed in the last 10 seconds. Between download attempts the algorithm waits a random time drawn independently from the uniform distribution on [.25, .75] sec.

In detail, bots base their decision on whether to initiate a download on two factors. One of these determinants is the variable "average delay"[3] (*AD*). The second factor is a configurable randomly drawn threshold value. In each period, bots (and real players in automatic mode) have three behavior settings that can be set by the experimenter. If they aren't defined for a given period, then the previous settings are used, and if they are never set, then the default settings are used. An example (using the default settings) is

AutoBehavior Player 1: MinThreshold 4000, RandomWidth 1000, PredictTrend Disabled

The definitions are:

1) MinThreshold (*MT*): The lowest possible threshold value in milliseconds. If the average delay is below this minimum threshold, then there is 100% certainty that the robot (or player in Auto mode) will attempt a download if not already downloading. The default setting is 4000 (= 4 seconds).
2) Random Width (*RW*): The random draw interval width in milliseconds. This is the maximum random value that can be added to the minimum threshold value to determine the actual threshold value instance. That is, $MT + RW = Max\ Threshold\ Value$.
3) Predict Trend (*PT*): The default setting is Disabled. However, when Enabled, the following linear trend prediction algorithm is used: $MT_2 = MT + AD_2 - AD$. A new Minimum Threshold (MT_2) is calculated and used instead of the original Minimum Threshold value (*MT*). The average delay (*AD*) from exactly 2 seconds ago (AD_2) is used to determine the new Minimum Threshold value.

A bot will attempt a download when $AD \le T = MT + RD$. A new threshold value (T) will be drawn (*RD* from a uniform distribution on [0, *RW*]) after each download attempt by the robot. Another important feature of the robot behavior is that a robot will never abort a download attempt.

To avoid artificial synchronization of robot download attempts, the robots check on *AD* every x seconds, where x is a uniformly distributed random variable on [.05, .15] seconds. Also, there is a delay (randomly picked from the uniform distribution on [.15, .45] seconds) after a download (successful or unsuccessful) has been completed and before the robot is permitted to download again. Both delays are drawn independently from each other and for each robot after each download attempt. The absolute maximum time a robot could wait after a download attempt ends and before initiating a new download (given that *AD* is sufficiently low) is thus 450ms + 150ms = 600ms.

APPENDIX B: STARCATCHER INSTRUCTIONS

UCSC 2/2003

I. GENERAL

You are about to participate in an experiment in the economics of interdependent decision-making. The National Science Foundation and other foundations have provided the funding for this project. If you follow these instructions carefully and make good decisions, you can earn a CONSIDERABLE AMOUNT OF MONEY, which will be PAID TO YOU IN CASH at the end of the experiment.

Your computer screen will display useful information regarding your payoffs and recent network congestion. Remember that the information on your computer screen is PRIVATE. In order to insure best results for yourself and accurate data for the experimenters, please do not communicate with the other participants at any point during the experiment. If you have any questions, or need assistance of any kind, raise your hand and somebody will come to you.

In the experiment you will interact with a group of other participants over a number of periods. Each period will last several minutes. In each period you earn "points" which are converted into cash at a pre-announced rate that is written on the board. You earn points by downloading stars. Each star successfully downloaded gives you 10 points, but waiting for a star to download incurs a cost. Every second that it takes to download the star will cost you 2 points. For example, if you start a download and it completes in 2 seconds, your delay cost is 4 = 2 points per second times 2 seconds. Therefore in this example you would earn 10 − 4 = 6 points.

Download delays range up to 10 seconds, depending on the number of other participants trying to download at the same time and background congestion. The delay cost can exceed the value of the download, so you can lose money when the network is congested. If the download takes 9 seconds you would earn 10 – 2*9 = –8 points, a negative payoff since the delay cost (18) is larger than the value of a star (10). Of course you can wait till the congestion clears: that way you don't make money, but neither will you lose any. Doing nothing earns you zero, but also costs zero.

II. ACTIONS

You have four action buttons: DOWNLOAD, RELOAD, STOP or GO TO AUTOMATIC. Clicking the DOWNLOAD button starts to download a star, and also starts to accumulate delay costs, until either:

- The star appears on your screen, so you earn 10 points minus the delay cost; or
- The star does not appear within 10 seconds, so you lose 20 points; or
- You click the STOP button before 10 seconds elapse, so you lose twice the number of seconds elapsed; or
- You click the RELOAD button. This is like hitting STOP and DOWNLOAD immediately after.

When you click GO TO AUTOMATIC a computer algorithm decides for you when to download. There sometimes are computer players (in addition to your fellow humans) who are always in AUTOMATIC. The algorithm mainly looks at the level of recent congestion and downloads when it is not too large.

III. SCREEN INFORMATION

Your screen gives you useful information to help you choose your action. The main window reports congestion on the network (how many people were downloading) in the last 10 seconds. The horizontal axis shows the delay time (from 0 to 10 seconds) and the height of each vertical bar represent the number of successful downloads. For example, in the 10 seconds slice of history shown in Figure 1, one successful hit took one second, 4 successful hits took two seconds, 10 took three seconds, 10 took four seconds, 4 took five seconds, etc. The color of the bar indicates whether the payoff from the download was positive (green) or negative (red). The Black bar on the right indicates the number of people who waited unsuccessfully for a star. The Blue bar (not shown in picture) indicates the number of people who hit Stop or Reload.

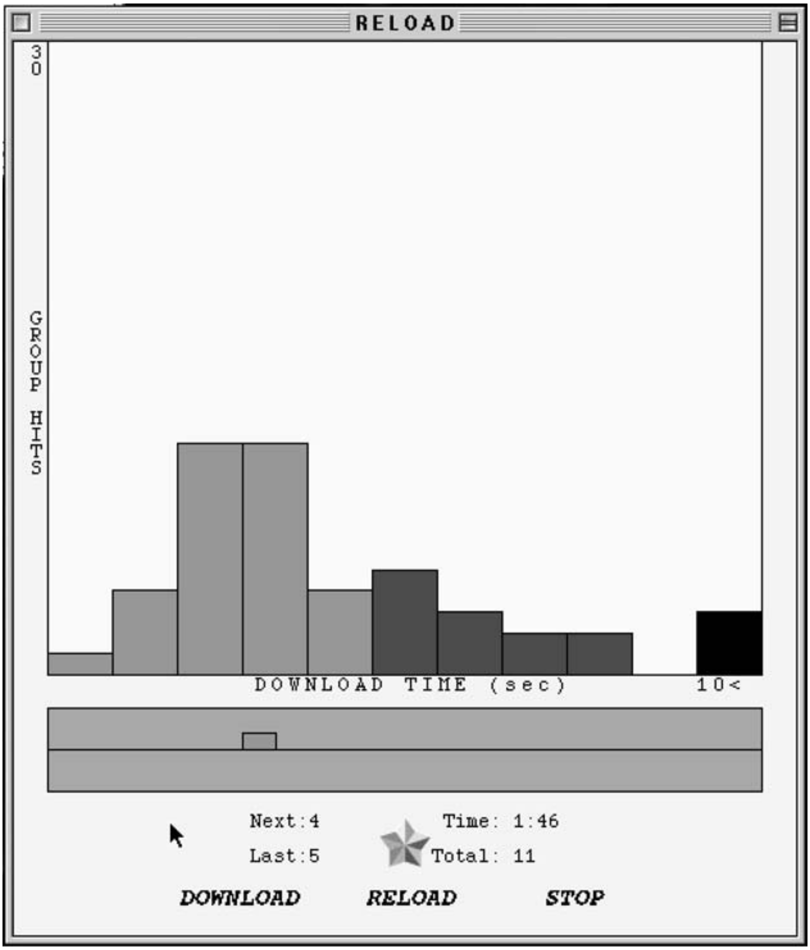

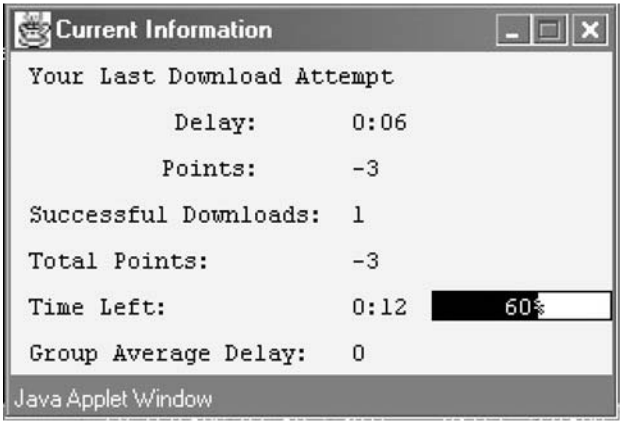

Just below the graph showing recent traffic is a horizontal status bar. This "status bar" has the same horizontal time scale as the graph above but shows the time of YOUR CURRENT download. When you click the "DOWNLOAD" button, a vertical bar will appear in the far left side of this status bar. The height of this bar represents the net payoff of a successful download *if it finished at that time*. As you wait for the download, this bar moves from left to right and shrinks as your delay costs accumulate. If the download takes so long that the delay cost exceeds the 10 pt. value of the star, this bar drops below the middle line, indicating a negative payoff.

NOTE: Pushing the STOP button at any point will give you a lower payoff than the bar indicates by 10 points since you will not get the value of the star but still pay the delay cost.

In the window "Current Information" you will find out how much time passed on your last download attempt (Delay), what your earnings were for the last download attempt (Points), the number of your successful downloads in this period (Successful Downloads), your total amount of points for this period (Point), the time left in the current period (Time Left), and the time needed for a download in the last 10 seconds, averaged across all players (Group Average Delay).

After the end of the first period two windows will appear on the right side of your screen. The top one displays information about your activity in the previous periods: number of attempted downloads (Tries), number of successful downloads (Hits), points (Winnings), your average points per try (Average), and a running total of your payoffs for all periods (Total). The bottom window shows the same statistics for the entire group. These windows will stay on your screen and will be updated at the end of each period.

IV. PAYMENT

The computer adds up your payoffs over all periods in the experiment. The last value in the 'Total' column in the 'Your Performance' window determines your

Your Performance

Period	Tries	Hits	Winnings	Average	Total
1	7	7	105	15	105
2	10	10	200	20	305

Warning: Applet Window

Group Performance

Period	Tries	Hits	Winnings	Average
1	40	31	158	4
2	32	27	119	4

payment at the end of the experiment. The money you will receive for each point will be announced and written on the board. After the experiment, the conductor will call you up individually to calculate your net earnings. You will sign a receipt and receive your cash payment of $5 for showing up, plus your net earnings.

V. FREQUENTLY ASKED QUESTIONS

Q: What happens if my net earnings are negative? Do I have to pay you?
A: No. To make sure that this never happens, you will be asked to leave the experiment if your total earnings start to become negative. In that case you would receive only the $5 show up fee.
Q: Is this some kind of psychology experiment with an agenda you haven't told us?
A: No. It is an economics experiment. If we do anything deceptive, or don't pay you cash as described, then you can complain to the campus Human Subjects Committee and we will be in serious trouble. These instructions are on the level and our interest is in seeing how people make decisions in certain situations.
Q: If I push STOP or RELOAD before a download is finished I get a negative payoff? Why?
A: Once you start a download, delay costs begin to accumulate. These costs are deducted from your total points even if you stop to download by clicking STOP or RELOAD.
Q: How is congestion determined?
A: Congestion is determined mainly by the number of download requests by you and other participants (humans and computer players). But there is also a random component so sometimes there is more or less background congestion.

Chapter 5

EXPERIMENTAL EVIDENCE ON THE ENDOGENOUS ENTRY OF BIDDERS IN INTERNET AUCTIONS

David H. Reiley[1]
University of Arizona

Abstract

This paper tests the empirical predictions of recent theories of the endogenous entry of bidders in auctions. Data come from a field experiment, involving sealed-bid auctions for collectible trading cards over the Internet. Manipulating the reserve prices in the auctions as an experimental treatment variable generates several results. First, observed participation behavior indicates that bidders consider their bid submission to be costly, and that bidder participation is indeed an endogenous decision. Second, the participation is more consistent with a mixed-strategy entry equilibrium than with a deterministic equilibrium. Third, the data reject the prediction that the profit-maximizing reserve price is greater than or equal to the auctioneer's salvage value for the good, showing instead that a zero reserve price provides higher expected profits in this case.

1. INTRODUCTION

The earliest theoretical models of auctions assumed a fixed number N of participating bidders, with the number commonly known to the auctioneer and the participating bidders. More recent models have relaxed this assumption, considering the possibility of costly bidder participation, so that the actual number of participating bidders is an endogenous variable in the model. In this paper, I use a field experiment, auctioning several hundred collectible trading cards in an existing market on the Internet, to test the assumptions and the predictions of models of auctions with endogenous entry.

I concentrate on three empirical questions in this paper. First, can an experiment turn up evidence of endogenous entry behavior in a real-world market? The answer to this question appears to be yes. Second, given the existence of endogenous entry, does the entry equilibrium appear to be better modeled as stochastic, or as deterministic? Evidence from the experiment indicates that the stochastic equilibrium concept is a better model of behavior. Third, is it possible to verify the theory of McAfee, Quan, and Vincent (2002, henceforth, MQV), that even with endogenous bidder entry, the optimal reserve price for the auctioneer to set is at least

A. Rapoport and R. Zwick (eds.), Experimental Business Research, Vol. II, 103–121.

as great as the auctioneer's salvage value? The answer to this question is "no," as a reserve price of zero appears to provide higher expected profits than a reserve price at the auctioneer's salvage value.

The field-experiment methodology of this study, that of auctioning real goods in a preexisting market, represents a hybrid between traditional laboratory experiments and traditional field research which takes the data as given. It shares with laboratory experiments the important advantage of allowing the researcher to control certain variables of interest, rather than leaving the researcher subject to the vagaries of the actual marketplace. (The key experimental treatment in this paper is the manipulation of the reserve price across auctions, to observe how participants react in their entry and bidding decisions.) It shares with traditional field research the advantage of studying agents' behavior in a real-world environment, rather than in a more artificial laboratory setting.

Although the experimental literature on auctions is vast,[2] almost all of these studies have imposed an exogenous number of bidders (determined by the experimenter). Three exceptions are Smith and Levin (2001), Palfrey and Pevnitskaya (2003), and Cox, Dinkin, and Swarthout (2001). Smith and Levin (2001) and Palfrey and Pevnitskaya (2003) design their experiments to determine whether the entry equilibrium which obtains is deterministic or stochastic, a question I also investigate in this paper. Cox, Dinkin, and Swarthout (2001) show that when participation in a common-value auction is costly, winner's-curse effects are attenuated.

In the empirical literature on auctions in the field,[3] one recent study considers endogenous entry. Bajari and Hortacsu (2003) note that in eBay auctions for coin proof sets, the number of observed bidders is positively correlated with the book value of the item and negatively correlated with the minimum bid for the item. From this they infer that bidding is costly, and they therefore provide a structural econometric model of bidding that includes an endogenous entry decision. The present paper adds to the empirical and experimental literatures on the endogenous entry of bidders by conducting a controlled experiment to gather evidence on the type of endogenous entry found in a real-world market.

The paper is organized as follows. The next section describes the relevant aspects of endogenous-entry auction theory, focusing on the testable implications. The third section describes the marketplace where the experiments took place, with twin subsections explaining the respective designs of the two sets of experiments. The fourth section presents the results, and a fifth section concludes.

2. THEORETICAL BACKGROUND

Recently, there have been a number of important extensions to Vickrey's (1961) original model of auctions with a fixed, known number of bidders. The earliest examples of endogenous-entry bidding models include Samuelson (1985), Engelbrecht-Wiggans (1987), and McAfee and McMillan (1987). In these models, bidders have some cost to participating (either the research required to learn one's value for the good, or the effort required to decide on a bid and submit it). This cost causes some potential bidders to stay out of the auction entirely, and can cause other effects as well. For example, Samuelson (1985) and Engelbrecht-Wiggans (1987), making different modelling assumptions, both find that endogenous entry drives down the auctioneer's optimal reserve price relative to a model of costless entry. One of the goals of the present paper is to demonstrate the existence of entry costs in a real-world auction market.

McAfee and McMillan (1987) model bidder entry as a pure-strategy, asymmetric Nash equilibrium. In these models, exactly n bidders enter the auction (out of a total of N > n potential bidders), and n is determined endogenously from the other parameters of the model (the auction format, the degree of affiliation of bidder values, the cost of entry, and so on). Alternatively, others have modeled a mixed-strategy, symmetric entry equilibrium (Engelbrecht-Wiggans (1987), Levin and Smith (1994), MQV). In the mixed-strategy models, bidders each enter with probability ρ, where ρ is determined endogenously.[4]

Levin and Smith (1994) point out that the difference between pure-strategy (deterministic) models and mixed-strategy (stochastic) ones has implications for social welfare: if entry is stochastic, then expected social surplus is decreasing in the number N of potential bidders. The reason is that the variance of the number n of actual entrants is increasing in N, and such variance is costly. In common-value auctions, then, it turns out that auctioneers can increase both social welfare and their own profits by using reserve prices to discourage entry.

In a separate paper, Smith and Levin (2001) perform an experiment in which they attempt to determine whether entry by bidders is stochastic or deterministic: they find evidence in favor of their stochastic hypothesis. However, the experimental procedure doesn't actually involve any auctions; rather, it assigns simulated auction payoffs by a lottery procedure.[5] Palfrey and Pevnitskaya (2003) modify this experimental design to conduct a first-price sealed-bid auction after the entry decision. They observe that the same bidders tend to enter repeated auctions, indicating a pure- rather than mixed-strategy equlibrium. Pevnitskaya (2004) provides a theoretical model of heterogenously risk-averse bidders to explain this observation. When some bidders are more risk-averse than others, and all bidders know this fact, the more risk-averse bidders stay out of the auction deterministically in order to collect a fixed payoff. Only the relatively less risk-averse bidders enter the auction, also deterministically.[6] Mixed-strategy equilibrium disappears in favor of a pure-strategy equilibrium the more risk-averse bidders stay out of the auction in favor of a fixed payoff, while relatively less risk-averse bidders enter the auction. In this paper, I attempt to provide evidence on the question of stochastic versus deterministic entry equilibria in a field environment.

MQV examine the effects of reserve prices where valuations are where bidder entry is endogenous and bidder valuations may be either affiliated. In their model, the auctioneer chooses a reserve price and announces her auction, together with the level of her reserve price, to N potential bidders. Bidders then decide whether or not to incur the participation costs, making a stochastic (mixed-strategy) entry decision. Next the participating bidders find out their private information about the value of the good, they submit their bids, and finally the auctioneer awards the good to the highest bidder. If no bidder chooses to enter and to bid at least the reserve price, then the auctioneer keeps the good for herself and earns some outside option utility, or "salvage value." The main prediction of MQV is that the optimal reserve price is at least as high as the salvage value of the good. This is a testable prediction; raising the reserve price from some lower value to the expected salvage value of the good should raise revenues for the auctioneer.

To summarize, this paper will attempt to answer three main questions. First, are entry costs relevant in the Internet auction market where I ran my experiments? Second, is the entry equilibrium a deterministic or a stochastic one? Third, is the optimal reserve price at least as high as the auctioneer's salvage value? Note that the first question is about an assumption of endogenous-entry, the second attempts to distinguish between two rival theories, and the third is a test of the empirical prediction of a specific model.

3. EXPERIMENTAL DESIGN

For this experiment, I auctioned trading cards via first-price, sealed-bid auctions, varying the reserve prices across treatments. The data in this paper are the same as in Lucking-Reiley (1999). The experiments took place in 1995 in a pre-eBay online market for collectible cards from *Magic: the Gathering*, a game which has enjoyed great success since its launch in August 1993. In the game, players assume the roles of dueling wizards, each with their own libraries of magic spells (represented by decks of cards) that may potentially be used against opponents. Cards are sold in random assortments, just like baseball cards, at retail stores ranging from small game and hobby shops to large chain retailers. The games's maker, Wizards of the Coast (now a division of Hasbro) has developed and printed thousands of distinct card types, each of which plays a slightly different role in the game.

As discussed in Lucking-Reiley (1999), soon after the introduction of *Magic*, players and collectors interested in buying, selling, and trading game cards began to use the Internet to find each other and carry out transactions. In a Usenet newsgroup dedicated to this purpose, traders used a variety of trading institutions, including negotiated trades of one card for another, sales at posted prices, and auctions of various formats, typically lasting multiple days.

Scarcity was one major determinant of transaction prices for cards, as some cards were printed in relatively low quantities, and some cards had gone out of print. The most common in-print cards were not worth trading over the Internet; their values were pennies or less. Cards designated "uncommon" but not "rare" traded for prices of ten cents to two dollars. Cards designated "rare" but still in print typically ranged in price from one to fifteen dollars. Out-of-print cards, depending on their initial scarcity and on other attributes, traded for as much as three hundred dollars. In this research project, I dealt only in out-of-print cards.

In addition to data generated in my own auctions, I also make use of contemporaneous market data from the weekly Cloister price list in this marketplace.

Cloister was a card trader who wrote a computer program that automatically searched the marketplace newsgroup for each instance of each card name (with some tolerance for misspellings) and gathered data on the prices posted next to each card name in the newsgroup messages. It then computed statistics for each card, and automatically archived these data on the Internet as a public service for other interested traders. Each card's reported list price is a trimmed mean over hundreds or thousands of different observations on the newsgroup. Despite some problems with these data, discussed in Lucking-Reiley (1999) many card traders adopted the Cloister price list as a standard measure of card market value, so I adopt it as a useful measure in my own analysis.

This marketplace represented an exciting opportunity to run auction field experiments. For the experiments, I purchased several thousand dollars' worth of cards (also via the Internet), and auctioned them off while systematically manipulating the reserve prices in order to observe their effects on bidder participation and bidding behavior. Because in any given week there were dozens of auctioneers holding *Magic* auctions on the Internet, as an experimenter I was able to be a "small player" who did not significantly perturb the overall market.

I employed two distinct experimental designs to collect the data. The first design examines the effects of a binary variable: whether or not minimum bids were used. By auctioning the same cards twice, once with and once without minimum bids, it exploits within-card variation to find the effects of the treatment variable on bidding and entry behavior. The second design investigates the effects of a continuous variable: the reserve price level (expressed as a fraction of the Cloister reference price). The between-card variation provides information that can be used to test the MQV prediction about the optimal reserve price level.

3.1. Within-Card Experiments

The first part of the data collection consisted of two pairs of auctions. Each of the four auctions was a sealed-bid, first-price auction of several dozen *Magic* cards auctioned off individually. This simultaneous auction of many different goods at once, although not common in other economic environments,[7] is the norm for auctions of *Magic* cards on the Internet. Running auctions in this simultaneous-auction format thus made the experiment as realistic and natural as possible for the bidders, who see many other similar auctions in the Internet marketplace for cards.

Each auction lasted for one week, from the time the auction was announced to the deadline by which all bids had to be received. I announced each auction to potential bidders via two channels. First, I posted three announcements to the appropriate Internet newsgroup. For each auction, I posted a total of three newsgroup messages spaced evenly over the course of the week of the auction. Second, I solicited some bidders directly via email messages to their personal electronic mailboxes. My mailing list for direct solicitation was comprised of people who had already demonstrated their interest in auctions for *Magic* cards by participation in previous ones.

The paired-auction experiment proceeded as follows. First, I held an absolute auction (no minimum bid) for 86 different cards (one of each card in the Antiquities expansion set). The subject line of the announcement read "Reiley's Auction #4: ANTIQUITIES, 5 Cent Minimum, Free Shipping!" so that potential bidders might be attracted by the unusually low minimum bid per card, essentially zero. (A 5-cent minimum is effectively no minimum, since the auction rules also required all bids to be in integer multiples of a nickel.) After the one-week deadline for submitting bids had passed, I computed the highest bid on each card. To each bidder who had won one or more cards, I mailed (electronically) a bill for the total amount owed.[8] After receiving a winner's payment via check or money order, I mailed them their cards. Almost no one defaulted on their winning bids.[9]

I also mailed a list of the winning bids to each bidder who had participated in the auction, whether or not they had won cards. This represented an effort to maintain my reputation as a credible auctioneer, demonstrating my truthfulness to those who had participated. I did not, however, give the bidders any explicit information about the number of people who had participated in the auction, or about the number of people who had received email invitations to participate.

After one additional week of buffer time after the end of the first auction, I ran the second auction in the paired experiment, this time with reasonably high minimum bid levels on each of the same 86 cards as before. The minimum bid levels were determined by consulting the standard (trimmed-mean) Cloister price list of *Magic* cards cited above of this paper, and setting the minimum bid level for each card equal to 90% of the value of that card from the price list.

This contrast in minimum bid levels (zero versus 90% of the Cloister price list) was the only economically significant difference between the two auctions.[10] By keeping all other conditions identical between the two auctions, I attempted to isolate the effects of minimum bids on potential bidders' behavior. One condition that could not be kept identical, unfortunately, was the time period during which the auction took place. Because the two auctions took place two weeks apart, there were potential differences between the auctions that might have affected bidder behavior. First, the demands for the cards (or the supplies by other auctioneers) might have changed systematically over time, which is a realistic possibility in such a fast-changing market as this one.[11] Second, since the auctions shared many of the same bidders, the results of the first auction may have affected the demand for the cards sold in the second auction.[12]

To control for such potential variations in conditions over time, I simultaneously ran the same experiment in reverse order, using a different sample of cards. This second pair of auctions each featured the 78 cards in the Arabian Nights expansion set, with minimum bids present in the first auction but absent in the second. Just as before, minimum bids were set at ninety percent of the market price level from the Cloister price list. The first auction in this pair began three days after the start of the first auction in the previous pair, so that the auctions in the two experiments overlapped in time but were offset by three days. Also, I used a larger mailing list for my email announcement in this pair of auctions (232 people) than I had for the previous pair of auctions (50 people), with the first mailing list being a subset of the second mailing list. Otherwise, all other conditions were identical between the two pairs of auctions.[13]

Table 1 shows a set of summary statistics for each of the four auctions in the within-card experiments.[14] The auctions are displayed in two pairs: first Auctions AA and AR, for the 86 Antiquities cards, and then Auctions BA and BR, for the 78 Arabian Nights cards.[15] Auctions AA and BA were with no minimum bids, while Auctions AR and BR had sizable minimums (equal to 90% of the market price).

The table contains quite a bit of descriptive information about the auctions, including the number of participating bidders, the number of bids received, and the total payments received from winning bidders. Note two key points. First, "real money" was involved in the auction transactions. Of the 73 different bills I sent to winning bidders over the course of the experiment, the median payment amount for each auction was between \$10 and \$24. A few individual payments even exceeded \$100.

Second, in each auction there are multiple winners. The number of winners in each auction ranges from 6 to 27, and the fraction of bidders who win at least one card is between 40 percent and 86 percent. In each auction, the median number of cards won by each winner is between 2 and 3.5, while the maximum number of cards won by a single bidder ranges from 12 to 26. Except in Auction AR, no winner won more than 29 percent of the cards sold in any single auction. (In Auction AR, participation was very low: only 7 people submitted bids, 6 of whom won at least one card, and 39 of the cards went unsold.) The biggest spender in any of the auctions won cards totalling \$316.50 of the total revenue of \$774.75 in Auction BA, generating 41 percent of the revenue despite winning no more than 15 percent of the cards – evidently, she was particularly interested in high-value cards. Thus, it is not the case that some people are the highest bidders on all cards in an auction, which suggests that a given bidder's valuations for different cards are at least somewhat independent. This gives some justification for reporting regression results in which each individual card bid is assumed to be an independent observation.

3.2. Between-Card Experiments

A second set of experiments was designed to examine the effects of changes in the *level* of the reserve price, rather than merely changes in the *existence* of reserve prices. Five first-price, sealed-bid auctions took place, each with a one-week timeframe for the submission of bids. Each was a simultaneous auction of many different items, this time with no overlap of items between auctions. Each card in the first four auctions (R1 through R4) had a posted reserve price. The fifth auction (R0) used a zero reserve price on every card, in order to provide a basis for comparison.[16] Just as before, I announced each auction via three posts to the relevant newsgroup, as well as via email to a list of bidders.[17]

In the first four auctions, I auctioned 99 different cards each time, setting a reserve price for each card as a particular fraction of the current Cloister price of that

Table 1. Summary statistics for within-card experiments.

	Auction AA	*Auction AR*	*Auction BA*	*Auction BR*
Minimum bids?	No	Yes	No	Yes
Card set	Antiquities	Antiquities	Arabian Nights	Arabian Nights
Start date	Fri, 24 Feb	Fri, 10 Mar	Tue, 14 Mar	Tue, 28 Feb
End date	Fri, 3 Mar	Fri, 17 Mar	Tue, 21 Mar	Tue, 7 Mar
Number of items for auction	86	86	78	78
Number of items sold	86	47	78	74
Revenue from twice-sold cards	$189.90	$234.75	$758.25	$783.80
Total auction revenue	$292.40	$234.75	$774.75	$783.80
Total number of bids	565	71	1583	238
Total number of bidders	19	7	63	42
from email invitations	12	5	44	35
from newsgroup announcements	7	2	19	7
Number of email invitations sent	52	50	232	234
Number of winners	15	6	25	27
Winner/bidder ratio	78.9%	85.7%	40.3%	64.3%
Cards per winner:				
Max	25	26	12	18
as share of total	29.1%	55.3%	15.4%	24.3%
Min	1	1	1	1
Mean	5.7	7.8	3.1	2.7
Median	3	3.5	2	2
Payment per winner:				
Max	$70.00	$129.40	$316.50	$128.00
as share of total	23.9%	55.1%	40.9%	16.3%
Min	$3.00	$0.70	$1.05	$2.55
Mean	$19.49	$39.13	$30.99	$29.03
Median	$10.50	$23.68	$13.15	$13.00

Table 2 Bids received in the within-card auctions.

	Auction AA	*Auction AR*	*Auction BA*	*Auction BR*
Minimum bids?	No	Yes	No	Yes
Card set	Antiquities	Antiquities	Arabian Nights	Arabian Nights
Number of bidders	19	7	62	42
Number of items for auction	86	86	78	78
Number of bids per bidder:				
Mean	29.7	10.1	25.5	5.7
Median	13.0	4.0	14.0	4.0
Max	86.0	29.0	78.0	30.0
Min	1.0	1.0	1.0	1.0

card. In each of the first two auctions, nine cards were auctioned at a minimum bid of 10 percent of the Cloister price, nine at 20 percent, nine at 30 percent, and so on, up to a maximum of 110 percent of the Cloister price. For each reserve-price level, I chose an assortment of different cards with widely different Cloister prices, and scattered the group randomly across the complete list of cards. After an analysis of the data from those auctions, I chose to collect more data both at very low and at very high reserve price levels. Therefore, the third and fourth auctions were designed to have equal numbers of cards auctioned at reserve levels of 10, 20, 30, 40, 50, 100, 110, 120, 130, 140, and 150 percent of the Cloister price.[18]

This variation in reserve price levels was designed to investigate how both bidder behavior and expected auction revenue would react to changes in the reserve price, and to calculate the optimal reserve price level. Normalizing by the Cloister price, since this is a standard reference price computed in the same way for all *Magic* cards, makes cross-card comparisons feasible. Besides the exceptions noted above, all experimental protocols and bidder instructions were kept identical to those used in the auctions with reserve prices in the experimental design described in section 3.1.

Summary statistics for the between-card auctions are given in Table 3. In auctions R1 to R4, reserve prices ranged from 0% to 150% of each individual card's Cloister value, and the average reserve price level varied slightly from auction to auction, from 60% to 85%. In auction R0, of course, the average reserve price level was zero.

As can be seen in the table, each auction had dozens of bidders and hundreds of bids on individual cards. The number of people receiving email invitations to

Table 3. Summary statistics for the between-card experiments.

	Auction R1	*Auction R2*	*Auction R3*	*Auction R4*	*Auction R0*
Card set	Artifacts	Black	White	Blue	Red/Green
Start date	Tue, 3 Oct	Fri, 6 Oct	Fri, 20 Oct	Mon, 23 Oct	Tue, 31 Oct
End date	Tue, 10 Oct	Fri, 13 Oct	Fri, 27 Oct	Mon, 30 Oct	Tue, 6 Nov
Number of items for auction	99	99	99	99	86
Number of items sold	98	92	77	78	86
Mean reserve level	60%	60%	85%	81%	0%
Total number of bids	798	652	366	401	1069
Total number of bidders	57	55	46	38	42
Number of email invitations sent	532	523	512	489	472
Total Cloister value	345.83	271.55	285.87	224.89	327.05
Total auction revenue	338.45	282.65	260.95	219.25	316.70
Revenue plus salvage	343.94	283.65	269.48	224.52	316.70

participate declined with each successive auction, but only due to recipients asking to be removed from my mailing list, so the changes in the mailing list should not have affected the number of potential participants. Note that the data from the between-card auctions is not directly comparable to that from the within-card auctions, because the size and composition of the pool of participating bidders changed considerably during the intervening six months. Very few bidders overlapped between the two experiments; most of the bidders in the between-card experiment were new recruits.

The table also displays aggregate statistics on revenue, including the total Cloister value of all the cards in each auction, the total revenue earned on cards which were sold, and a grand-total revenue figure which also includes the salvage value of the unsold cards. The auction revenue in each case was reasonably close to the total Cloister value of the cards; in Auction R2 I earned revenue greater than the total Cloister value, while in the three others I earned slightly less.

4. RESULTS

I now present the results from the experiment, separately for each of the three empirical questions outlined above. Are entry costs relevant? Is the entry equilibrium stochastic or deterministic? Do the auctioneer's profits improve as he raises the reserve price to be at least as high as his salvage value?

4.1 Entry costs are relevant

The within-card experiments demonstrate that endogenous bidder entry appears to be the right model for this market. Statistics on the number of card bids per participating bidder are shown in Table 2. As expected, individual bidders tend to submit fewer bids in the presence of minimums than they do in the absence of minimums. This does not in itself demonstrate the existence of bidding costs; a bidder who contemplates how much to bid and then decides that the reserve price exceeds his maximum willingness could still be counted as having "participated," because the decision cost would already have been incurred even though the reserve price prevents me from observing a low bid. In the auctions with minimums, no single bidder submitted bids on even half of the cards; the maximum number of bids by a single bidder was 30. By contrast, there were bidders in both of the no-minimum auctions who submitted individual bids on every single card.

Interestingly, relatively few bidders followed this strategy of bidding on every single card in the absolute (no-minimum) auctions. Only one out of 19 bidders bid on every single item in Auction AA, and only six of 62 bidders bid on every single item in Auction BA. These statistics indicate that the cost of submitting a bid (the participation cost) is high enough to affect bidder behavior, and thus this experimental environment is appropriate for exploring endogenous-entry bidding models such as MQV. If there were no cost to submitting a bid, then one would expect to see all of the participating bidders submitting bids on every card (as low as a nickel, say), since every card does have some positive resale value even to people who get no consumption utility from it.[19] I conclude that bidders deem the probability of getting a bargain (and thus a resale profit) on such a card is low enough that the expected profit from bidding does not always outweigh the cost of having to decide on a bid amount and to type the approximately ten characters required to submit another card bid. Indeed, the median number of card bids submitted by a single bidder was only 13 (of a possible 87) in Auction AA, and 14 (of a possible 78) in Auction BA, even though these auctions had no minimum bids.

Thus, bidders do appear to make a participation decision consistent with the existence of small entry costs; the number of participating bidders in each auction is not exogenous. The classical theory makes some accurate predictions about the effects of reserve prices, as shown earlier, despite this violation of its assumptions.

4.2. Is the entry equilibrium stochastic or deterministic?

Given the existence of endogenous entry, I now ask: is the entry equilibrium deterministic or stochastic? Very few bidders bid on a card both times it was offered, despite the fact that the same people were invited each time. Nineteen and seven bidders, respectively, bid in the two Antiquities auctions, but only 4 people overlapped between the two auctions. In the Arabian Nights auctions, there were 42 and 62 bidders, but only 17 of the bidders overlapped between the two. Thus, in each pair of auctions, there were a proportionally large number people who entered the first auction but not the second, and other people who entered the second auction but not the first. This argues in favor of a stochastic equilibrium, as the most natural kind of deterministic equilibrium is one in which the same bidders enter each time.

Two objections might be raised to the result just presented. First, it might be the case that people enter one auction but not the other because the latter auction has reserve prices which are higher than they are willing to pay. However, this screening-out explanation cannot account for the bidders who bid in the presence of reserve prices but fail to bid in the absence of reserve prices; there were 3 such bidders in the Antiquities auctions, and 25 such bidders in the Arabian Nights auctions. The second potential objection is that bidders may have bid in the chronologically first auction, but not the second, in a pair because they had already bought the cards by the time the second auction occurred. This complaint potentially affects the 25 Arabian Nights bidders just cited, who bid in Auction BR but not in Auction BA. Indeed, three of these bidders each placed a bid on a single card in Auction BR and won it, so there would be no reason to expect them to bid in the second auction. However, none of the remaining 22 bidders won all the cards they bid on in Auction BR: ten did not win any cards at all, while the remaining twelve won an average of 50 percent of the cards they bid on. It is still possible that these bidders managed to purchase the rest of the cards they were interested in from someone else during the week that passed between my two auctions, but I can at least say that they did not buy them from me. Thus, the evidence is fairly strong that bidders in these auctions made stochastic entry decisions: faced with the same auction opportunity, the same person might sometimes enter and sometimes fail to enter. This contrasts with the results of Palfrey and Pevnitskaya (2003), who find evidence that the less risk-averse bidders consistently tend to enter while the more risk-averse bidders consistently tend not to enter.

The stochastic entry decision might not be due to conscious randomization by a bidder trying to follow a “mixed strategy” in the textbook sense. Perhaps bidders enter “randomly” because of other things happening in their lives: a college student had too much homework one week, or a computer programmer had a family emergency. Lots of random events could cause bidders to show up to one auction but not another. However, in terms of auction design and welfare considerations, what matters is whether the entry decisions in a real-world auction are deterministically predictable by the auctioneer and by the rival bidders. My evidence shows that at least in this market, bidder entry decisions are stochastic, so the model of Levin and Smith (1994) has empirical relevance.

4.3. Optimal Reserve Price with Endogenous Entry

Recall that the main prediction of the MQV paper is that raising the reserve price from some low value to the salvage value of the good will increase expected auction profits, even in an endogenous-entry context. In order to understand the effect of the reserve price on expected revenues, I turn to the between-card experimental data. Recall that these data provide samples of auction revenues at differing reserve price levels (normalized by Cloister price for each card).

Table 4 summarizes the results of the experiment separately for each reserve-price decile, from reserve prices of 0% of the Cloister price to reserve prices of 150% of the Cloister price. The table displays the total number of cards I auctioned at each reserve price, the number of those which went unsold, and the mean and standard deviation of the revenues at each reserve price level. The revenues are also normalized by the Cloister price of each card, and an unsold card counts as an observation of zero revenue. The data are displayed graphically in Figure 1, with the mean revenues plotted against the reserve prices. The error bars show one standard error in each direction (where the standard error equals the standard deviation in revenues for that reserve price level divided by the square root of the number of observations at that reserve price level). We see that the revenues are quite high at a reserve price of zero, then drop off sharply at a reserve price of 10% of Cloister price. Revenues seem to rise again, generally, between 50% and 100% of the Cloister price, then fall again at higher reserve price levels. There are surprisingly high revenues observed at 140% to 150% of the Cloister price, albeit with high standard errors.

To test the MQV prediction also requires an estimate of the salvage value for the unsold cards. I asked my local card dealer what he would pay me for my unsold cards; he responded with an offer that was 20 percent of their Cloister price. He further indicated that 20% of Cloister price would be his average offer price for

Table 4. Cards and revenues at each reserve price. (Reserve prices and revenues normalized by the Cloister price of each card).

Reserve price	*Total cards*	*Unsold cards*	*Mean revenue*	*Std dev of revenue*
0.0	96	0	1.192	1.071
0.1	33	0	0.847	0.549
0.2	36	2	0.857	0.594
0.3	34	0	0.823	0.599
0.4	27	2	0.775	0.500
0.5	32	2	0.945	0.517
0.6	20	1	0.977	0.480
0.7	16	0	0.965	0.259
0.8	23	1	1.093	0.469
0.9	31	4	0.983	0.500
1.0	35	6	1.055	0.674
1.1	32	4	1.113	0.491
1.2	15	7	0.804	0.861
1.3	21	10	0.760	0.772
1.4	17	6	0.967	0.867
1.5	14	5	1.104	0.893

cards of this quality and quantity, so I adopt a salvage value of 20% percent of the Cloister price for each card.[20]

Now the question is whether a reserve price equal to the salvage value yields expected profits at least as high as a reserve price less than the salvage value (0% or 10%) of salvage value. The point estimates of revenues certainly indicate that the opposite is the case. In order to perform a formal hypothesis test, first I calculate expected profits rather than expected revenues. For the 0% reserve price, all cards sold, so profits remain the same as revenues: 1.192. For the 20% reserve price, two cards went unsold; when I count salvage profits of 20% of Cloister price for each of

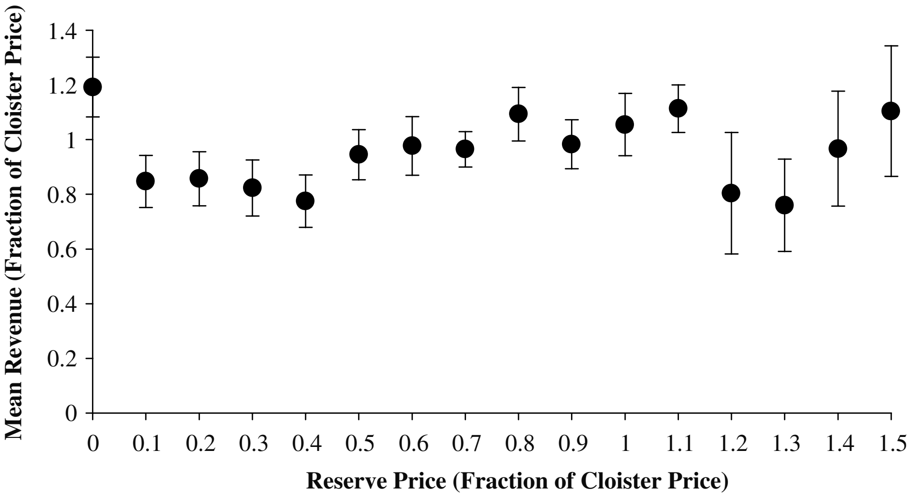

Figure 1. Mean Revenue as a Function of Reserve Price

these cards, the estimate of expected profit rises slightly, from 0.857 to 0.870. Using the calculated standard deviations, I conduct a test of the null hypothesis of equality between expected profits at 0% reserve price and expected profits at 20% reserve price. The resulting standard normal test statistic is 2.18, with a p-value of 0.029. Thus, I reject the null hypothesis of equality at the 5% level of significance, and conclude that expected profits are actually *higher* for a zero reserve price than they are for a reserve price equal to the salvage value.[21]

This is a violation of the theoretical prediction, an example of a case in which the auctioneer does better to hold an auction with a zero reserve price than to set the reserve price equal to the salvage value. One possible explanation is that an auction with no reserve price generates more enthusiasm among bidders, causing higher levels of participation. In other words, modest minimum bids may eliminate some high valuation-bidders, who would have bid high if they had participated, but decide not to participate unless their attention is attracted by an auction with zero minimum bids. Although a few items may end up being sold at very low prices, they might serve as "loss leaders," similar to the goods advertised at deep discounts by supermarkets, enabling the auctioneer to collect higher revenues overall. This proposed effect involves increased entry through attracting bidders' *attention*, with the absolute auction as a type of promotion, rather than assuming the bidders will make a careful calculation of the costs versus the benefits of bidding. Note in Table 4 that the total number of bidders in Auction R0 is actually lower than in the other auctions, which might seem to be evidence against this effect, although I should also note that the number of cards in auction R0 is also lower than in the auctions with reserve prices. One caveat about this finding is that most of the zero-reserve-price cards were sold in the same auction (R0). Although I did attempt to keep all other variables constant across auctions, the anomaly might be due to some uncontrolled factor which was different between R0 and the earlier auctions.

5. CONCLUSIONS

This study presents the results of controlled experimental auctions performed in a field environment. By auctioning real goods in a preexisting, natural auction market, I have obtained data in a manner that is intermediate between laboratory experiments and traditional studies of field data. Some variables were unfortunately unobservable and uncontrolled – for example, I could not assign "valuations" for each good to each bidder, as a laboratory experimentalist might. On the other hand, I have the opportunity to hold constant most of the relevant variables in the environment, and to manipulate the treatment variable, which in this case was the existence and level of reserve prices. By giving up the ability to observe and manipulate some of variables that laboratory experimenters can control, I gained a realistic environment. The participants had previous experience bidding for the types of real goods I auctioned, and the auctions took place in an Internet-based market where bidder entry decisions seemed potentially important.

The first result is that entry costs are an important feature of this real-world auction markets, thus confirming the central assumption of endogenous-entry auction theory. The costs in the *Magic*-card market are probably not nearly as dramatic as those postulated in other markets (for example, in the market for offshore oil rights the bidders typically hire geologists to perform extensive analysis of the potential for oil in a given tract). Here, the cost of acquiring information about individual cards is quite small, but even the cost of typing in a bid amount appears to have observable effects.

Second, when the same cards were auctioned twice in rapid succession, very different sets of people decided to submit bids, despite the fact that the same superset of people were invited to participate both times. This can be interpreted as evidence in favor of the stochastic (mixed strategy) entry equilibrium model, where the number of participating bidders varies unpredictably.

Third, I found that, contrary to the theory of McAfee, Quan, and Vincent (2002), a zero reserve price can earn higher expected profits than a reserve price equal to the auctioneer's salvage value. Perhaps an absolute auction attracts significantly more bidder attention than an auction with even modest reserve prices, causing more additional entry than might be suggested by a model of rationally calculated bidder entry decisions. It will be interesting to see whether this finding can be replicated in other auction markets.

NOTES

[1] Department of Economics, the University of Arizona. I wish to thank Mike Urbancic, Marius Hauser, and Mary Lucking for their research assistance, and Skaff Elias for product information about *Magic: the Gathering*. I would like to thank J.S. Butler, Rachel Croson, Glenn Ellison, Elton Hinshaw, Dan

Levin, Kip King, Preston McAfee, Rob Porter, and Jennifer Reinganum for advice and constructive criticism.

[2] See Kagel (1995) for a review of auction experiments.

[3] See Hendricks and Paarsch (1995) for a review of empirical work on auctions.

[4] These models find simple, symmetric solutions by assuming that bidders decide whether to participate before they learn their valuations. In my auctions, it is reasonable to assume that participants had information about their valuations before making the entry decision, so the entry outcome might be asymmetric. An example of such an asymmetric model is given by Samuelson (1985), where only those bidders with high valuations participate in the auction, and the entry equilibrium is in pure strategies.

[5] Subjects made the decision whether or not to incur the cost c to enter. After the entry outcome was observed, each of the n entrants had a 1/n chance of winning the payoff for that round of the experiment.

[6] An interesting implication of Pevnitskaya's model is that an auctioneer can actually make himself worse off by advertising a sealed-bid auction heavily. An increase in the number of potential bidders increases the self-selection effect, causing less and less risk-averse bidders to enter the auction, and thereby causing less aggressive bidding, as risk aversion is well-known to increase bids in a first-price sealed-bid auction.

[7] Although simultaneous auctions are not traditional for familiar auctions, such as those of art, estate goods, or tulip bulbs, such formats have been used for timber and offshore oil auctions. The advent of computerized bidding appears to be making the simultaneous auction format even more common. In addition to the card auctions in this newsgroup market, simultaneous Web-based auctions are becoming common at commercial sites such as eBay, and a simultaneous-auction format was used for the recent FCC auctions of spectrum rights (see McMillan (1994) for details).

[8] Although the standard practice in this marketplace is for auctioneers and other card sellers to charge buyers for postage and/or handling, I chose not to do this. I wanted bidders to bid independently, as much as possible, on each of the cards in which they were interested. Someone seriously interested in one card might decide to bid higher on a second card in the same auction than they would if the cards were auctioned independently, because they would like to spread out the postage costs per card by purchasing more than one card simultaneously from the same source. In addition, some of the cards I auctioned had rather low values, and I wanted to avoid having the card values be swamped by the cost of shipping. For example, if a bidder won a single card for 20 cents and then had to pay a fixed 50-cent shipping charge on top of that, the amount of useful information which could be derived from her bid would be rather suspect. Therefore, in the interests of keeping bid data as clean as possible, I decided to pay postage costs myself, and announced in advance that first-class shipping was included in the amount of each bid.

[9] A small number of winning bidders failed to pay for the cards they had won. In all, I received payment for 90% of the cards sold, constituting 89% of the reported revenue in the within-card auctions. Almost all of the "deadbeat" bidders were those who won only a single card, and who explained that they had originally hoped to win more cards, and didn't feel it was worth it to complete the transaction. I discouraged such behavior, but was unable to eliminate it. Only one or two individuals won multiple cards but failed to pay for them. Since none of the unpaid cards seemed to have outlandishly high winning bid amounts, I have taken the point of view that all bids were made in good faith, and have not excluded any observations from my analysis.

[10] Both auctions lasted exactly seven days. The same 86 cards were up for bid in each auction. Each auction announcement was posted exactly three times to the marketplace newsgroup, and was emailed to primarily the same list of potential bidders. Even the subject line of the announcements and mailings was kept identical, except that in the second auction, the words "5 Cent Minimum" were removed.

[11] For example, certain cards from the Arabian Nights expansion set increased in value by a factor of ten during their first year out of print. It turns out that market prices for cards were actually rather stable during the month in which this experiment was conducted, but I did not know *a priori* what was going to happen to card prices.

[12] For example, suppose that a particular bidder is anxious to obtain a single Guardian Beast card for her deck, so that her valuation of the card is higher than that of any of the other bidders in the experiment. She may win the card in the first auction, and then have zero demand for that same card in the second auction. If this is generally the case for most cards, that the highest-value bidders in the sample are screened out in the first auction, then we might expect to see systematically lower revenues in the second auction.

[13] A sample auction announcement, as it appeared to the potential bidders both in electronic mail and in the market place newsgroup, can be found on the World Wide Web at: http://eller.arizona.edu/~reiley/papers/EndogenousEntry.html

[14] A note on mnemonics. The first letter represents the card set: A for Antiquities, B for Arabian Nights. The second letter is A for an absolute auction (reserve prices equal to zero), and R for an auction with positive reserve prices.

These auctions were part of a series of auctions run for a larger research program, so participating bidders saw me run several other auctions (not part of the research presented here) during the same time period. This had two advantages where the experimental design is concerned. First, it helped avoid drawing bidders' attention to the point of my research. (For example, during this period I also ran an English auction and a second-price auction and another first-price auction, with different sets of cards.) I feared that if they knew I was looking for the effects of reserve prices, it might distort their behavior (for example, they might consciously try to bid consistently from one auction to another). Second, it had the effect of making bidders unsure of what I would do next. In particular, I didn't want bidders to expect that I would always auction the same card twice, for it might distort their behavior if they knew they would have a second chance to bid on the same card.

[15] A few of the auction items I denote as "cards" were actually groups of cards: either a sealed packet of out-of-print cards, or a set of common cards bundled together.

[16] It was necessary to do another absolute auction, rather than just reusing those of the previous section, because those took place in a substantially different time period, with a very different number of invited bidders, thus making their results incomparable to those of the within-card experiments.

[17] For this series of auctions, the bidder pool was quite a bit larger than before. 531 individuals were emailed to participate in Auction R1, and as some people specifically requested to be removed from my auction announcement mailing list, the list dropped to 489 individuals by the time Auction R4 began.

[18] In practice, the number of cards at each reserve-price level ended up not being precisely equal. Because I required bids to be in multiples of $0.05, I always took the computed reserve price and rounded it down to the nearest acceptable bid amount. In the analysis below, I take the ratio of the actual reserve price used to the Cloister price, and round to the nearest 10% level in order to examine the effects of the reserve price on expected revenue. This results in unequal numbers of cards at the different levels of reserve prices.

[19] Because of the time and transaction costs involved in selling it off, it is conceivable that for some bidders, the net resale value of a card might be less than five cents. However, most cards had gross resale values of over a dollar, and many bidders in this market could be assumed to take some pleasure in trading cards with others, as trading is a big part of the game culture.

[20] I might have been able to shop around for a better price with a different card dealer, but this represents my best estimate of a salvage value, which by definition should be net of all administrative costs, including search costs.

[21] The reader might wonder about robustness to alternative assumptions about the salvage value. In particular is possible that I may have overstated auction profits, because my revenue figures are not discounted for the labor and postage I spent in order to ship the cards to the winning bidders, and therefore non-auction salvage might actually be more attractive, relative to auction revenues, than I initially assumed. Assuming salvage values of 30%, 40%, or even 50% of Cloister price still yields statistically significantly higher profits for a zero reserve price than for a reserve price equal to salvage value, so the result is quite robust. (The difference is no longer statistically significant for assumed salvage values of 60% or higher, as there are fewer observations at these higher reserve price levels.)

REFERENCES

Bajari, Patrick and Ali Hortacsu. (Summer 2003). "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions." *RAND Journal of Economics*, 34(2), 329–355.

Black, Jason. *Cloister's Magic Card Price List*, <http://www.hhhh.org/cloister/ pricelists/>, various weekly issues.

Cox, James C., Samuel H. Dinkin, and James T. Swarthout. (October 2001). "Endogenous Entry and Exit in Common Value Auctions." *Experimental Economics*, 4(2), 163–181.

Cox, James C., Bruce Roberson, and Vernon L. Smith. (1982). "Theory and Behavior of Single Object Auctions," in *Research in Experimental Economics*, Vernon L. Smith, ed., Greenwich, Conn.: JAI Press.

Engelbrecht-Wiggans, Richard. (June 1987). "On Optimal Reservation Prices in Auctions." *Management Science*, 33(6), 763–770.

Engelbrecht-Wiggans, Richard. (1992). "Optimal Auctions Revisited." *Games and Economic Behavior*, 5, 227–239.

Hendricks, Kenneth, and Harry J. Paarsch. (1995). "A Survey of Recent Empirical Work Concerning Auctions." *Canadian Journal of Economics*, 28(2), 315–338.

Kagel, John H. (1995). "Auctions: A Survey of Experimental Research," in *The Handbook of Experimental Economics*, J. Kagel and A. Roth, eds. Princeton: Princeton University Press, 501–585.

Levin, Dan, and James L. Smith. (1994). "Equilibrium in Auctions with Entry." *American Economic Review*, 84(3), 585–599.

Levin, Dan, and James L. Smith. (1996). "Optimal Reservation Prices in Auctions." *Economic Journal*, 106, 1271–1282.

Lucking-Reiley, David. (1999). "Using Field Experiments to Test Equivalance Between Auction Formets: Magic on the Internet." *American Economic Review*, 89(5), 1063–1080.

McAfee, R. Preston, and John McMillan. (1987a). "Auctions and Bidding." *Journal of Economic Literature*, 25(2), 699–738.

McAfee, R. Preston, and John McMillan. (1987b). "Auctions with Entry." *Economics Letters*, 23(4), 343–347.

McAfee, R. Preston, Daniel Quan, and Daniel Vincent. (2002). "How to Set Optimal Minimum Bids, with an Application to Real Estate Auctions." *Journal of Industrial Economics*, (50)4, 391–416.

McMillan, John, (1994). "Selling Spectrum Rights." *Journal of Economic Perspectives*, 8(3), 145–162.

Milgrom, Paul R., and Robert J. Weber. (1982). "A Theory of Auctions and Competitive Bidding." *Econometrica*, 50, 1089–1122.

Palfrey, Thomas R., and Svetlana Pevnitskaya, (August 2003). "Endogenous Entry and Self-Selection in Private-Value Auctions: An Experimental Study." Caltech Social Science working paper #1172.

Pevnitskaya, Svetlana, (January 2004). "Endogenous Entry in First-Price Private-Value Auctions: The Selection Effect." Ohio State University working paper.

Reiley, David H. (August 2004). "Field Experiments on the Effects of Reserve Prices in Auctions: More Magic on the Internet," University of Arizona working paper.

Riley, John G., and William Samuelson, (1981). "Optimal Auctions." *American Economic Review*, 71(3), 381–392.

Samuelson, William F. (1985). "Competitive Bidding with Entry Costs." *Economics Letters*, 17, 53–57.

Smith, James L, and Dan Levin, (2001). "Entry Coordination, Market Thickness, and Social Welfare." International Journal of Game Theory, 30, 321–350.

Vickrey, William. (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance*, 16(1), 8–37.

Wilson, Robert. (1992). "Strategic Analysis of Auctions," in *The Handbook of Game Theory*, R.J. Aumann and S. Hart, eds. New York: North-Holland, 227–279.

Chapter 6

HARD AND SOFT CLOSES: A FIELD EXPERIMENT ON AUCTION CLOSING RULES

Daniel Houser
George Mason University

John Wooders
University of Arizona

Abstract

Late bidding in online auctions has attracted substantial theoretical and empirical attention. This paper reports the results of a controlled field experiment on late bidding behavior. Pairs of $50 gift certificates were auctioned simultaneously on Yahoo! Auctions, using a randomized paired comparison design. Yahoo? site allows sellers to specify whether they wish to use a hard or soft close, and this enabled us to run one auction in each pair with a soft close, and the other with a hard close. An advantage to our randomized paired design is that differences in numbers of bidders, numbers of simultaneously occurring auctions and other sources of noise in bidding behavior are substantially controlled when drawing inferences with respect to treatment effects. We find that auctions with soft-closes yield economically and statistically significantly higher mean seller revenue than hard-close auctions, and that the difference is due to those cases where the soft-close auction is extended.

1. INTRODUCTION

The growth of auctions on the Internet raises new theoretical questions, provides a wealth of data on bidding behavior in auctions, and presents new opportunities for running experiments in the field. The present paper reports the results of a field experiment on the effects of closing rules on auction outcomes. Different auction sites have adopted different closing rules. On eBay, auctions have a "hard" close, with the seller specifying when it ends (either exactly 3, 5, 7, or 10 days after it is listed). On Amazon, auctions have a "soft" close, with the auction ending at the scheduled closing time if no bids arrive in the prior 10 minutes, but with the auction otherwise ending only after 10 minutes has elapsed without a bid.

The present study takes advantage of the fact that Yahoo! Auctions allows a seller, when listing an auction, to choose whether to end the auction with a hard or

A. Rapoport and R. Zwick (eds.), Experimental Business Research, Vol. II, 123–131.

soft close.[1] In our experiment, identical pairs of $50 gift cards were auctioned simultaneously, with one card of the pair auctioned with a soft close and the other auctioned with a hard close. We find that soft-close auctions yield higher revenue than hard-close auctions, and this difference is statistically significant. Both types of auctions were equally likely to have a "late bid", i.e., a bid submitted within the last five minutes of the auction. However, our ability to detect differences in the frequency of late bidding is limited by the small sample size of our study.

Our study is motivated, in part, by Roth and Ockenfels' (2000, 2002) comparison of last minute bidding (also know as "sniping") on eBay and Amazon, on auctions of computers and antiques. Roth and Ockenfels find that there is significantly more late bidding on eBay auctions than on Amazon auctions. In their data set, more than two-thirds of the eBay auctions received a bid in the last 30 minutes of the auction, and about 40 percent received bids in the last five minutes. In contrast, on Amazon only about one quarter of the auctions received a bid in the last 30 minutes of the auction, and only 3 percent received a bid in the last five minutes.

This difference in the timing of bids is consistent with a theoretical analysis of hard and soft close auctions. One explanation for the difference stems from the fact that, in practice, there is some chance that an attempt to place a bid at the last minute of an auction will not be successful. When this is this case, Roth and Ockenfels (2000) show that for auctions with a hard close there is an equilibrium in which all bidders submit a bid equal to their value at the last minute (under some assumptions on the distribution of values). In this equilibrium the bidders tacitly collude – all the bidders respond to an early bid by bidding their values immediately. In equilibrium a bidder prefers to bid late, and face a smaller number of competing bids, rather than bid early and having his bid successfully placed, but face competing bids from all the other bidders. Roth and Ockenfels also show that last-minute bidding is a best response to an "incremental bidding" strategy by naïve bidders.[2] In soft-close auctions, a last-minute bid extends the bidding. Roth and Ockenfels show that in soft-close auctions it is not an equilibrium for all bidders to submit last-minute bids. Nor is last-minute bidding a best response to incremental bidding in soft close auctions.[3]

Both theoretical explanations of late bidding suggest that seller revenue is lower in auctions with a hard close. In the equilibrium with tacit collusion the seller receives (in expectation) fewer bids. Against an incremental bidder, a bidder who snipes pays less than the incremental bidder's value.

Several factors prevented Roth and Ockenfels from comparing seller revenue in hard and soft close auction. When their data was collected in the fall of 1999, eBay was already the dominant auction venue, with many more bidders than Amazon.[4] Even if the same items were sold on both sites, this alone would make it difficult to determine whether revenue differences between hard and soft close auctions were due to differences in the closing rule or in the number of bidders. In fact, the computers and antiques sold on each auction sites are heterogeneous both within and across the auction sites. The sellers on the two sites also have different reputations (represented by their feedback profiles), which influences the bidders' values for the items.[5] These factors prevent a straightforward comparison of revenues of hard-close (eBay) and soft-close (Amazon) auctions.

Our experiment had a paired design, with pairs of identical items auctioned at the same time (on Yahoo), with one item in the pair sold in a soft-close auction and the other sold in a hard-close auction. Hence the number of potential bidders and their characteristics were identical for both auctions in a pair. The same seller ID was used for both auctions, and hence the seller's feedback profile (called the seller "rating" on Yahoo) was also the same between paired auctions. This design allows for a test of the effect of the closing rule on revenue, and it has high power with even a small sample of auctions. The results of the present paper support the conclusion that revenue is lower in hard-close auctions.

2. RELATED EXPERIMENTAL LITERATURE

Several other papers have also investigated the effect of the closing rule on the timing of bids and seller revenue. We focus on the results for seller revenue. In a laboratory experiment, Ariely, Ockenfels, and Roth (2002) find that seller revenue is higher in the soft-close treatment than in the two hard-close treatments they consider. (In one hard-close treatment, last minute bids are processed with probability .8, while in the other they are processed for sure.) The soft-close also yields more revenue than a second-price sealed-bid auction.

In a paper closely related to our own, Gupta (2001) studies the effect of closing rules by comparing the outcomes of hard and soft-close Yahoo auctions. His approach involved selling forty matched pairs of identical sealed music CD's, with one CD from each pair being sold in an auction of each type. He found that the mean sale price in the soft-close auctions was \$6.89, as compared to \$6.60 in the hard-close auctions. However, he reports that this price difference is not statistically significant ($p = 0.31$). More generally, he found that "comparisons between the two treatment groups [hard and soft-close auctions] yielded no significant differences in either price, bid number or bid timing" (p. 26).

Gupta's study was carefully done. Nevertheless, one potentially important reason that he did not find differences in behavior between auction types is that the participants in his auctions might not have realized that they were bidding in a hard- or soft-close auction, and even if they recognized it, might not have understood the meaning of the closing rule. Evidence in support of this is that although several of his auctions were extended, none of his extended auctions received bids during the extended time. In the present study, we avoid this confound by making salient on our auction page the nature and meaning of the auction closing rule (see the Item Information in Figure 1). Another possible explanation for the difference between our results and Gupta's is that the stakes in his study are substantially smaller, and hence may not provide bidders with sufficient incentive to carefully time the placing of their bids.

Moreover, although Gupta auctioned matched pairs of items, it is not clear whether he auctioned each item in the pair concurrently. Final auction prices can

Figure 1. Typical Soft-Close Auction Page.

vary for a large number of reasons, particularly because of variations in the number of potential bidders. As a result, the impact of closing rules can be obscured by other differences in the auction environment. As we describe in detail below, our design is to run each item in the pair at the same time, and therefore ensure a common auction environment. This reduces the effect of confounding factors on outcome differences and, consequently, allows relatively more compelling inference about closing rule effects.

3. EXPERIMENT DESIGN

Our design provides a clean and simple way to compare the effects of different closing rules on auction outcomes. The primary advantage of our field experiment is that we gain a subject pool and environment more closely tied to the naturally occurring world. At the same time, we inevitably lose some control that we have in the laboratory. Like all empirical analyses of field auction data, we lose control of the number of potential bidders (a number critical for the theory), as not all potential bidders are observable. (The number of actual bidders is of course observable, but this provides only a lower bound on number of potential bidders.) We also lose control over all dimensions of the set of competing auctions, including how many there are, how closely related they are to our product, and how they are advertised.

Our approach to mitigating the noise associated with the field experiment to adopt a randomized "paired" experimental design. The idea is to run two auctions simultaneously, where the auctions are identical in every way except the closing rule. In particular, one of the auctions is listed with a hard close, and one with a soft close. The main advantage to this randomized paired design is that differences in numbers of bidders, numbers of simultaneously occurring auctions and other sources of noise in bidding behavior are substantially controlled when drawing inferences with respect to closing rule effects. We chose Yahoo because Yahoo allows sellers to specify whether they want to use a hard or soft close.

One potential disadvantage of the paired design is that our auctions compete with each other, and some might argue that this creates an artificial environment that weakens our study's external validity. In fact, a casual inspection of any major auction website reveals many essentially identical auction listings across many product categories. Our experience is that it is more the exception than the rule to have a unique item with few very closely competing auction listings. Consequently, although it forces a departure from some of the premises of standard auction theory, we believe a paired design enhances our study's ability to predict the effects of different closing rules as used in actual Internet auctions.

The item sold in each of our auctions was a $50 gift certificate that could be redeemed at a chain-store with outlets throughout the United States. Although each pair of auctions sold a gift certificate for the same store, the stores were varied across auction pairs. The stores were chosen in an effort to appeal to customers of varying demographic characteristics, so that we would obtain variety in the people interested in participating in our auctions. For example, we auctioned gift certificates to both Sears and Crabtree and Evylyn. While certainly there is some overlap in these stores' customers, this overlap is not likely perfect. The seven stores we included in our study are: Borders Books, Circuit City, Crabtree and Evylyn, Sears, Target, Toys-R-Us, and Victoria's Secret.

An important advantage of selling gift certificates, then, is that they allow high homogeneity within a pair yet provide heterogeneity across pairs. There are other substantial advantages to selling gift certificates. An important one is that gift certificate auctions are clearly *private value* auctions. That is, one bidder's bid does not convey any information to the other bidders about the value of the gift certificate. For example, a bidder's value for a Borders Books certificate will depend on idiosyncratic factors including his cost of traveling to the nearest Borders, and his preference for Borders products in relation to those available at other nearby bookstores. This latter could vary with, for example, relative return policies. Other practical advantages to selling gift certificates are that they are easy to obtain, easy and inexpensive to ship, easy to describe and, again, exceptionally homogenous.

Both auctions in a pair were posted at the same time and using a nearly identical page layout and item description. Figure 1 shows the auction page for a typical soft-close auction. The text "Auction may get automatically extended," which appeared in the page's "Notes" section, indicated to participants that the auction had a soft close. A hard-close auction contained, instead, the text "This auction does not get extended automatically." In addition, we described the closing rule for each type of auction in the item description. Soft-close auctions included the text "This auction is automatically extended an additional 5 minutes whenever a bid is placed within 5 minutes before the auction close," whereas in hard-close auctions we stated "This auction does not get automatically extended and ends at the close time given above." As discussed above, the reason for emphasizing the closing rule was to increase the likelihood that subjects would both notice and understand this auction feature. Note again that, other than differences regarding the closing rule, the auction pages were identical.

An undergraduate research assistant created a Yahoo account for the purpose of this project and posted each auction pair. The account was held fixed across auctions. All auction winners were promptly and appropriately sent the item they had won. As a result, the seller's rating score increased over the course of the experiment. This is not a concern for our study, as each auction in a pair was held in the same reputation environment, and our inferences are based on the distribution of within-pair outcome differences.

4. RESULTS

We conducted 15 pairs of auctions during the Fall 2001 academic semester. One auction pair was lost due to a recording error (a Victoria's Secret auction) leaving 14 auction pairs in our data set. While this number is not large, it should be remembered that we base our results on differences in auction outcomes within pairs, a procedure that has relatively high statistical power. Indeed, we see below that even with this limited data set, statistical differences in outcomes between auctions with hard and soft closes are apparent.

Table 1 describes the outcomes of the 14 auctions in our data set. The first column lists the store associated with the auctioned \$50 certificates. Note that five of our seven stores were used for two auction pairs, Sears was used for one and Borders for three. The next three columns describe the number of bids, revenue and whether there were late bids in each of the soft-close auctions. The number of bids ranged from a low of six (Sears) to a high of 33 (Toys-R-Us) with a mean of 18. Revenues varied between \$27.25 (Borders) and \$46.05 (Target), with an average of \$36.15. Five of the soft-close auctions received late bids and were extended. Within this set, the number of bids ranged from 15 to 33, while revenues ranged from \$27.25 to a maximum \$35.33.

The next three columns of Table 1 detail the results of the hard-close auctions. The number of bids ranged from a low of 5 (Circuit City) to a high of 37 (Victoria's Secret), with a mean of just under 19. Revenue from the hard-close auctions was lowest in a Border's Books auction (\$26) and highest in a Target auction (\$47), averaging just under \$35. There were five hard-close auctions in which bids were entered within 5 minutes before the close. (Of course, these auctions were not

Table 1. Auction Outcomes

	Soft Close			*Hard Close*			*Difference (Soft – Hard)*	
Store	*Number of Bids*	*Revenue*	*Late Bids*	*Number of Bids*	*Revenue*	*Late Bids*	*Number of Bids*	*Revenue*
Borders	17	$37.01	No	12	$35.01	No	5	2.00
Target	9	$46.00	No	19	$47.00	No	–10	–1.00
Victorias Secret	29	$34.00	Yes	37	$31.95	Yes	–8	2.05
Sears	6	$31.01	No	13	$32.01	No	–7	–1.00
Toys “R” Us	33	$35.33	Yes	32	$33.00	No	1	2.33
Circuit City	18	$34.33	No	5	$31.01	No	13	3.32
Crabtree & Evelyn	21	$32.06	Yes	27	$28.03	Yes	–6	4.03
Borders	21	$32.01	No	19	$32.01	No	2	0.00
Target	16	$46.05	No	23	$45.00	No	–7	1.05
Toys R Us	12	$41.00	No	15	$41.00	No	–3	0.00
Victorias Secret	10	$42.00	No	16	$43.00	No	–6	–1.00
Circuit City	28	$37.00	No	18	$38.21	Yes	10	–1.21
Borders	17	$27.25	Yes	12	$26.00	Yes	5	1.25
Crabtree & Evelyn	15	$31.11	Yes	17	$26.04	Yes	–2	5.07
Mean	18.00	$36.15		18.93	$34.95		–0.93	1.21

extended.) Among the late bid set, the number of bids ranged from 12 to 37, and revenues from $26 to $38.21. Four of the 14 hard-close auctions generated revenues greater than $40.

The price of the same gift card varied substantially across auctions at different times. Borders cards, for example, fetched as much as $37.01 in one soft-close auction, but received only $27.25 in another soft-close auction. This suggests that the Yahoo gift card market is relatively “thin,” with the price depending heavily on the willingness to pay of the bidders who happen to be present. (Note that prices in the hard-close auctions are correlated with prices in the soft-close auctions.) This variation in prices highlights the advantage of the paired design. It controls for the substantial variation in price that is due to factors other than the auction closing rule.

The final two columns of Table 1 detail the differences in outcomes between the soft and hard-close auctions. The difference reported is outcome in the soft-close auction less the outcome in the hard-close auction. With respect to number of bids, this difference ranges from a low of −10 (Target) to a high of 13 (Circuit City). The average difference is about −1, but is not statistically significant. The implication is that the number of bids in the two environments is about the same. Revenue differences range from a high of about \$5 (Crabtree and Evelyn) to a low of −\$1.21 (Circuit City). There were two occasions in which the auction types earned identical revenue (Borders and Toys-R-Us.) In eight of 14 of our auction pairs the soft-close auction earned more revenue. The soft-close auctions generated an average (over all auctions) of \$1.21 (3.5%) more than the hard-close auctions, and this difference is statistically significant (Wilcoxon signed-rank test for paired observations, $p < 0.05$).

A closer inspection of the revenue difference figures reveals a very close relationship of revenue to whether the soft-close auction was extended. In particular, on each of the five occasions where the soft close auction received late bids, it also generated higher revenue than the hard-close auction. Among this set, the average revenue advantage was about \$3 (about 10%). The soft-close auction earned greater revenue in only three of the nine auctions that did not include late bidding, and among that set the mean revenues were almost exactly identical.

In summary, our results suggest that soft-close auctions produce statistically significantly greater revenue on average than hard close-auctions, but this is due to those cases where the auction is extended. An interesting feature of our data is that there are an equal number of late bids placed in each type of auction.

5. CONCLUSION

Laboratory evidence from Ariely, Ockenfels, and Roth shows that a seller obtains more revenue when they sell using a soft rather than a hard-close auction. This study presents evidence that the soft-close auction continues to be superior, even when it is employed in the field. Furthermore, the soft-close auction raises more revenue than a hard-close auction, even when both auctions must compete for bidders, as is the case in the field.

The difference between our results and Gupta’s (2001) suggests that the size of the stakes may be important in understanding behavior in soft- and hard-close auctions. In particular, the revenue advantage we find for soft-close auctions may become insignificant in auctions of smaller denomination gift cards, if bidders believe is it not worth their effort to time the placing of their bids. This is an interesting direction for future research. A larger field study, using more auctions than the present study, would provide more insight into whether the closing rule affects the timing of bids.

ACKNOWLEDGMENT

We gratefully acknowledge financial support from the SRP Initiative on Technology, Public Policy and Markets (Univ. of Arizona) and the International Foundation for Research in Experimental Economics. Part of this work was completed while Wooders was a visitor at Hong Kong University of Science and Technology. He is grateful for their hospitality.

NOTES

[1] The closing rules are slightly different between Amazon auctions and Yahoo soft-close auctions. A Yahoo soft-close auction ends at the scheduled closing time if there are no bids in the 5 minutes prior to the close. Otherwise, the auction is extended by 5 minute increments, until one of these increments passes without any bids. Hence, while an Amazon auction may end any number of minutes after the scheduled close, a Yahoo soft-close auction always ends a multiple of 5 minutes after the scheduled close.

[2] An incremental bidder raises his bid by the minimum increment whenever he is outbid, so long as this would not lead him to bid above his value.

[3] See also Ariely, Ockenfels, and Roth (2002) for theoretical models of late bidding in eBay and Amazon auctions. In common value auctions they show that an expert bidder, who is better informed about the item's true value, also has an incentive to bid late so that other bidders can not free ride on his information.

[4] Bidders may also self select into eBay or Amazon auctions in a way that depends on their characteristics, introducing the possibility of selection bias.

[5] In a study of Pentium processor auctions on eBay, Houser and Wooders (2000) show that positive and negative feedback both have a statistically significant effect on price.

REFERENCES

Ariely, Dan, Ockenfels, Axel, and Roth, Alvin. (December 2002). "An Experimental Analysis of Ending Rules on Internet Auctions," mimeo, Harvard University.

Gupta, Neeraj. (2001). "Internet Auctions: A Comparative Study of Seller Options on eBay, Amazon, and Yahoo!" Undergraduate thesis, Harvard College.

Houser, Dan and Wooders, John. (2000). "Reputation in Auctions: Theory, and Evidence from eBay." University of Arizona Working paper #00-01.

Roth, Alvin and Ockenfels, Axel. (2000). "Last Minute Bidding and the Rules for Ending Second-Price Auctions: Theory and Evidence from a Natural Experiment on the Internet." National Bureau of Economic Research Working paper No. 7729.

Roth, Alvin and Ockenfels, Axel. (Month 2001). "Last Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet." *American Economic Review*, volume, 1341–78.

Chapter 7

WHEN DOES AN INCENTIVE FOR FREE RIDING PROMOTE RATIONAL BIDDING?

James C. Cox
University of Arizona

Stephen C. Hayne
Colorado State University

Abstract

Economics has focused on models of individual rational agents. But many important decisions are made by small groups such as families, management teams, boards of directors, central bank boards, juries, appellate courts, and committees of various types. For example, bid amounts in common value auctions such as the Outer Continental Shelf oil lease auction are typically decided by committees. Previous experimental research with natural groups has found that group bidders are significantly less rational than individual bidders in how they use information in common value auctions. Experiments reported here involve cooperative and non-cooperative nominal groups. The unequal profit-sharing rule applied to non-cooperative nominal groups creates an incentive to free ride within the bidding groups. This incentive to free ride tends to offset the winner's curse and promote rational bidding.

1. INTRODUCTION

Economics has traditionally focused primarily on the behavior of individual rational agents interacting in markets and other strategic game environments. But many important economic, political, scientific, cultural, and military decisions are made by groups. Decision-making groups have many forms including families, management teams, boards of directors, central bank boards, juries, appellate courts, and committees of various types.

Numerous researchers in management science and psychology have previously studied group decision-making. Our research involves some important departures from previous work in that: (a) we study group decision-making in the context of strategic market games, rather than non-market games against nature; and (b) we use a natural quantitative measure to determine whether and, indeed, how far groups' decisions depart from rationality.

A. Rapoport and R. Zwick (eds.), Experimental Business Research, Vol. II, 133–149.

We study group decision-making in the context of bidding in common value auctions. Bidding strategies in many important auctions are usually decided by groups. For example, oil companies typically use committees comprised of managers and geologists to determine bids for purchasing oil leases (Capen, Clapp & Campbell, 1971; Hoffman, Marsden, & Saidi, 1991). General contractors typically use committees to determine bids for large contracts (Dyer and Kagel, 1996).

In another paper (Cox and Hayne, 2002), we study decisions made by individuals and by "natural" groups – groups whose members conduct face to face discussion to arrive at a single group decision by whatever decision rule they choose to adopt. In this paper, we study decisions made by "nominal" groups – groups whose members arrive at a group decision by some pre-specified decision rule without an opportunity for face-to-face discussion. Nominal groups are further divided into "cooperative" groups (where there is no conflict of interest among group members), and "non-cooperative" groups (where the interests of the individual group members are partially conflicting).

Decision-making responsibility may be assigned to groups, rather than individuals, because of a belief that (a) groups are inherently more rational than individual decision-makers and/or (b) important pieces of information are possessed by different individual members of groups. In Cox and Hayne (2002), we report some perhaps surprising results comparing bids made by individuals with bids made by natural groups of 5 individuals that share equally in the profit or loss from a winning bid. The question posed in that paper is whether natural groups are more or less rational than individuals in common value auctions. We report that the answer depends upon the defining characteristics of natural groups. If one assumes that natural groups are decision-making entities consisting of more than one individual with distinct information then comparison of results from treatments involving natural groups, with value signal sample size of 5, with treatments involving individuals, with signal sample size of 1, supports the conclusion that natural groups are less rational than individuals. On the other hand, if one assumes that natural groups consist of individuals that have common information then comparison of results from treatments involving natural groups, with signal sample size of 1, with treatments involving individuals, with signal sample size of 1, supports the conclusion that natural groups are neither less nor more rational than individuals.

In the present paper, we compare bidding behavior of cooperative nominal groups with that of non-cooperative nominal groups; we change the incentives within the group. The treatments in this experiment involve two categories of groups with three individuals in each group. Bidding occurs under two conditions that differ only with regard to the way in which the group's profit or loss from a winning bid is divided among the group members. This enables us to vary the relationship within the group while keeping intact the number of decision-makers in each group and the nature of their joint decision vis-a-vis the other bidding groups in the market. For both types of nominal groups, the imposed decision rule is that a group's bid is the average of the bids submitted by individual members of the group. In the cooperative group treatment, all members of a group share equally in the profit or loss from a winning bid. In the non-cooperative treatment, the individual members share equally in the value of the item if their group has the winning bid but share unequally in the cost of the winning bid: each member of the winning-bid group pays one-third of the amount of his own bid. While all the members in a non-cooperative group have a common interest in submitting a winning bid, each individual member also has an interest in minimizing his own cost associated with the bid. Thus the non-cooperative treatment provides each player with an incentive to free ride. Consider individual j who is a member of a bidding group. Individual j prefers that her group-mates submit bids that are high enough for her group to have the winning bid while she bids zero. But if individual j bids too low she runs the risk that her group will not have the high bid, and hence receive zero payoff.

While the non-cooperative treatment involves an incentive to free riding, whether or not a low bidder actually free rides depends on the bids of the other group members. If the low bidder prevents his group from winning an auction in which, because of their high bids, the other two group members would have lost money, then he is not gaining at their expense, hence not actually free riding. In contrast, if a group has the high bid in spite of a low bid by one member of the group, then the low bidder is free riding. In any case, the non-cooperative treatment does provide an incentive to group members to free ride on others' bids that are high enough to win the auction.

The behavior of cooperative nominal groups can be compared with that of the natural groups studied in Cox and Hayne (2002). This comparison allows one to separate the effect on group decision-making of processes associated with face-to-face interaction from the effect of the mere aggregation of individuals' decisions. In contrast, comparison of the bidding behavior of cooperative and non-cooperative nominal groups allows one to isolate the effect on group bidding behavior of the incentive to free-ride created by the unequal sharing rule for non-cooperative groups. This permits us to address the question of when, or indeed if an incentive for free riding promotes rational bidding in common value auctions.

2. A QUANTITATIVE MEASURE OF DEVIATION FROM RATIONAL BIDDING BEHAVIOR

Consider an auction market where the bidders do not know the value of the item being sold when they submit their bids and the value, v of the item is the same for all bidders. Consider a first-price sealed-bid auction in which the high bidder wins and pays the amount of its bid for the auctioned item. Further envision that each bidder receives an independent signal, s_i that provides an unbiased estimate of the object's true value. The expected value of the auctioned item conditional on the bidder's signal is denoted by $E(v|s_i)$. The expected value of the auctioned item conditional on the bidder's own signal being the highest of N signals (i.e., equal to the highest order statistic, y_N) is denoted by $E(v|s_i = y_N)$. For bids by $N > 1$ bidding entities, one has

$$E(v|s_i) = s_i > E(v|s_i = y_N) \tag{1}$$

by well-known properties of order statistics. Thus if bidders naively submit bids equal to or slightly below their common value estimates, s_i they will have an expected loss from winning the auction; that is, they will suffer from the winner's curse.

Consider the case where the common value of the auctioned item is uniformly distributed on the interval, $[v_l, v_k]$ and each individual agent's signal is independently drawn from the uniform distribution on $[v - \theta, v + \theta]$. For this case, one has

$$E(v \mid s_i = y_N) - E(v \mid s_i) = -\frac{N-1}{N+1}\theta, \tag{2}$$

for all $s_i \in [v_l + \theta, v_h - \theta]$. Thus, if bidders naively submit bids equal to their signals then the cursed winning bidder will have an expected loss in the amount $\theta(N - 1)/(N + 1)$. This expected loss is increasing in the number of bidders, N.

Now assume that *each member* of each group has a signal that is independently drawn from the uniform distribution on $[v - \theta, v + \theta]$. Under these conditions, groups of size $G > 1$ have a signal sample size of $G > 1$ on which to base their estimates of the common value. Because the signals are drawn from a uniform distribution, the signal sample midrange, m_i provides an unbiased estimate of the value of the auctioned item. The expected value of the auctioned item conditional on the bidder's signal sample midrange is denoted by $E(v \mid m_i)$. The expected value of the auctioned item conditional on the bidder's own signal sample midrange being the highest of N signal sample midranges (i.e., equal to the highest order statistic of sample midranges, z_N) depends on the sample's range, r_i and is denoted by $E(v \mid r_i, m_i = z_N)$. With signals drawn from the uniform distribution, one has

$$E(v \mid r_i, m_i = z_N) - E(v \mid m_i) = -\frac{N-1}{N+1}\left(\theta - \frac{1}{2}r_i\right). \tag{3}$$

Comparison of statements (2) and (3) reveals that groups with size $G > 1$ signal samples will have a smaller expected loss from naively bidding their signal sample midranges, than will individuals from naively bidding their signals, except in the improbable extreme outcome in which all of the signals in the sample with the highest midrange have the same value (and, hence $r_i = 0$). In the other improbable extreme outcome, in which the signal sample with the highest midrange has a range equal to 2θ, there will be zero expected loss from bidding an amount equal to the sample midrange. But, of course, in this case the bidder knows the auctioned item's value with certainty.

Note that equation (3) suggests a quantitative criterion for determining the extent of deviation from minimally-rational bidding in common value auctions. The magnitude of the winner's curse that is exhibited by winning bidder i, in a market with N bidders, is

$$EVCurse = b_i^w - E(v \mid r_i^w, m_i^w = z_N) \tag{4}$$

where b_i^w is the winning bid, v is the common value of the auctioned item, r_i^w is the winning bidder's signal sample range (which is zero for signal sample size of 1), m_i^w is the winning bidder's signal sample midrange (or signal, for signal sample size 1), and z_N is the *Nth* order statistic of sample midranges (or signals, for signal sample size 1). Note that *EVCurse* is the magnitude of the expected loss (or profit, if it is negative) from winning the auction.

In order not to have an expected loss from winning, a bidding group or individual must discount its naive estimate of the common value (its signal or its signal sample midrange) by at least the amount $(\theta - \frac{1}{2}r_i)(N - 1)/(N + 1)$, where it is understood that $r_i = 0$ for signal samples of size 1. Furthermore, the size of this minimum rational discount is independent of m_i so long as $m_i \in [v_l + \theta, v_h - \theta]$. These conditions are essentially always satisfied by the signals drawn in our experiments in which $\theta = 1800$ and $[v_l, v_k] = [2500, 22500]$. Therefore, deviations from minimally-rational bidding can be measured by linear regressions relating winning bids to the signal sample ranges and midranges of the winning bidders, as we do in section 4. Bids that yield zero expected profit are given by the following equation when $m_i \in [v_l + \theta, v_h - \theta]$:

$$b^{Zero} = -\frac{N-1}{N+1}\theta + m_i + \frac{N-1}{2(N+1)}r_i. \tag{5}$$

Bids that are lower than b^{Zero} have non-negative expected profit regardless of what rival bidders are bidding. In contrast, bids that are higher than b^{Zero} are not economically rational because they have non-positive expected profit, that is, they exhibit the winner's curse. Of course, a bid may be less than b^{Zero} but still higher than the Bayesian-Nash equilibrium bid for the bidder's signal sample. But that would not imply that the bid is irrational unless the bidder knew that all other bidders were bidding according to the Bayesian-Nash equilibrium bid function. If rival bidders are bidding above the Bayesian-Nash equilibrium bid function then a bidder's rational best reply may be to also bid higher than bids specified by the equilibrium bid function. But it would never be rational to bid higher than b^{Zero} because: (a) such bids have non-positive expected profit regardless of what rival bidders are bidding; and (b) such bids have negative expected profit if rival bidders are bidding less than b^{Zero}. In short, a bidder who bids above b^{Zero} will have negative expected profits unless there are "bigger fools" to save him from having the high bid by making bids with negative expected profits themselves. Therefore, it is comparison with b^{Zero} that provides a strong test for rational bidding. Any bid above b^{Zero} is characterized by the winner's curse, which is not rational bidding.

3. EXPERIMENTAL DESIGN AND PROCEDURES

The design of our experiment addresses the effects on nominal groups of different profit-sharing rules. Individuals are randomly assigned to three-person groups. Each market includes three nominal group bidders. A nominal group's bid is the average (that is, the mean) of the three bids submitted by individual members of the group. In the cooperative-group treatment, the members of the group with the highest average bid share equally in the profit or loss from submitting the winning bid. In the non-cooperative-group treatment, each member of the highest-bid group receives one-third of the value of the auctioned item and pays one-third of the amount of his individual bid.[1] This unequal profit-sharing rule produces an incentive for free riding within groups: an individual member of a group will realize the highest possible profit (conditional on the common value) when the bids by the other two group members are high enough to win the auction and she bids zero. But if a subject bids too low he runs the risk that his group will not have the high bid and incur an opportunity cost of foregone profit.

Reports of results from previous common value auction experiments with individual bidders (Kagel and Levin, 1986; Kagel, et al., 1989) have focused on the behavior of experienced subjects, where "experience" means having participated in one or more previous common value auction experiments. The reason for this is that most subjects fall victim to the winner's curse in *all* experimental treatments when they are first-time bidders but such inexperienced behavior is not considered to be very interesting. We use subject experience as a treatment to allow comparison of our results with those in the literature.

The experiment was conducted in the Economic Science Laboratory (ESL) at the University of Arizona. Each individual had his own personal computer that was connected via the Internet to software running on a server at Colorado State University. Subjects were recruited from the undergraduate student population. Experimental sessions were run in two-day sequences of two-hour blocks. Subjects were paid all of their earnings at the end of the experiment on the second day. Subjects were randomly assigned to groups and randomly dispersed to computers as they moved into the laboratory.

The subjects were given written instructions describing bidding procedures in the first-price sealed-bid auctions. The instructions are reproduced in the appendix. The instructions contain a detailed description of the information environment of the common value auctions. Thus, subjects were informed in non-technical terms that in each auction round the computer would draw a value for the auctioned item from the discrete uniform distribution on the integers greater than or equal to 2,500 experimental dollars and less than or equal to 22,500 experimental dollars. They were informed that the common value would not be revealed but that it would be the midpoint of a uniform distribution from which their value estimates, or signals, would be independently drawn. They were informed in non-technical terms that, after the computer drew a common value v for a round, it would draw all signals independently from the uniform distribution on $[v - 1800, v + 1800]$. Information about how the signals would be drawn was presented to the subjects both in their written instructions and orally by the experimenters. The oral presentation used the analogy with bidding on oil leases and interpreted the signal(s) as estimates of the value of an oil lease by geologist(s). The instructions did not contain any discussion of the order statistic property that is conventionally thought to underlie the winner's curse. The instructions contained non-technical explanations of how the common values and subjects' signals were generated, the rules of the first-price sealed-bid auction, and the applicable profit-sharing rule. The subjects were informed orally of the number of bidding groups in the auction and the number of subjects in each group. This information was also written prominently on a whiteboard at the front of the laboratory.

On day 1, the inexperienced subjects first participated in 10 periods of practice auctions. After each practice auction, the subjects' computer monitors displayed the common value, all subjects' bids, and the amount won or lost by the high bidder. The subjects were each given a capital endowment of 1,000 experimental dollars in order to allow them to make at least one sizable overbid without becoming bankrupt. At the end of the practice rounds, the subjects' profits and losses were set to zero and they began at least 30 monetary-payoff rounds with new 1,000 experimental dollar capital endowments. The actual number of monetary-payoff rounds to be completed was not announced. Signals were presented to the subjects on sheets of paper; each subject was given a single sheet of paper with signals for 10 practice rounds and 40 monetary-payoff rounds. The experiment was ended on a monetary payoff round randomly chosen between 30 and 40. Signals, common values, and bids were denominated in experimental dollars, with a clearly specified exchange rate into U.S. dollars. During the monetary-payoff rounds the information reported at the end of each auction included only the common value and the high bid, not the bids by other bidders. We decided not to report all bids in order to make collusion more difficult and to adopt procedures that correspond to minimal reporting requirements in non-laboratory auctions. Each individual in a winning bidder nominal group could see his individual profit or loss and cumulative balance. The procedures were the same on day 2 as on day 1. Earnings from both day 1 and 2 sessions, together with the \$15 individual participation fees, were paid after the end of the day 2 sessions.

A few individuals made data entry errors during the experiment sessions with *inexperienced* subjects even though the software asked them to confirm their bids. Such errors were obvious because they usually consisted of mistakenly typing one fewer or one more digit in the bid than was intended. Subjects who made these errors usually immediately brought them to the experimenters' attention. Such errors were usually obvious because they produced bids that were too low or too high by a multiple of 10. We forgave losses resulting from data entry errors. Auction period data with data entry errors are excluded from our data analysis for the inexperienced bidders. There were no known data entry errors in the experiment sessions with *once-experienced* subjects.

Many *inexperienced* group-bidding entities made winning bids that turned out to be so high that they attained negative cumulative payoffs. A few once-experienced groups also incurred negative balances in the cooperative treatment but none did so in the non-cooperative treatment (see Table 1). When a group's cumulative payoffs were negative at the end of a session, the loss was forgiven (the group was permitted to "go bankrupt"). Allowing bidders to continue bidding after they have attained a negative cumulative balance can be a problem because the experimenter might lose control of their incentives. Therefore, we analyze data in Section 4.2 only from

Table 1. Summary Bidding Behavior

Experience	*Groups*				*Payoff*[a]		
	Type	*Total*	*Bankrupt*	*Zero Bids*	*Average*	*Low*	*High*
0	Coop.	36	19	12	\$6.31 (\$18.93)	\$0 (\$0)	\$23.09 (\$69.26)
0	Non-coop.	24	2	7	\$12.92 (\$38.77)	\$0 (\$1.60)	\$42.99 (\$87.61)
1	Coop.	36	9	12	\$11.25 (\$33.74)	\$0 (\$0)	\$25.40 (\$76.21)
1	Non-coop.	24	0	1	\$23.01 (\$69.03)	\$0 (\$22.75)	\$77.39 (\$103.05)

a. Figures in parentheses are earnings for the whole group.

periods in a bidding market session prior to a period in which any bidding group had a negative cumulative balance at the beginning of the period.

4. DATA ANALYSIS

4.1. Group Payoffs and Bankruptcies

The nominal groups' and individual subjects' high, low, and average money payoffs from bidding in all rounds of all sessions in the common value auctions (excluding the participation fees) are reported in Table 1.[2] The first column reports the experience level of the groups (inexperienced = 0, once-experienced = 1) and the second column shows the experimental treatment (Cooperative or Non-cooperative). The individual subject payoff amounts are, by definition, equal to one-third of the group amounts except for the low payoff and high payoff amounts for the non-cooperative treatment. There were large differences between the lowest and highest payoffs in all treatments. The average payoff in non-cooperative sessions was about twice what it was in cooperative sessions for both inexperienced and once-experienced subjects. This reflects the much higher incidence of the winner's curse, resulting in many more bankruptcies in the cooperative treatment than in the non-cooperative treatment. More than half of the inexperienced groups in the cooperative treatment went bankrupt while only two inexperienced groups went bankrupt in the non-cooperative treatment. The rate of bankruptcy decreases with subjects' experience in both treatments but remains very different. Nine out of 36 once-experienced groups became bankrupt in the cooperative treatment while none did so in the non-cooperative treatment. Thus the incentive to free ride, by submitting low bids, in the non-cooperative treatment increases average profits from bidding and decreases bankruptcies resulting from the winner's curse. However, it is interesting to note that the incidence of zero bids is *not* higher in the non-cooperative treatment than in the cooperative treatment.[3]

Further insight into bidding behavior in the cooperative and non-cooperative treatments is provided by analysis of data from periods in which no groups were bankrupt and hence, as explained in section 3, there is not a concern that the experimenters may have lost control of some subjects' incentives. Data analysis reported in Tables 2 and 3 excludes data from all market periods after any bidding group attained a negative cumulative balance. This has a large effect on data used for inexperienced subjects, especially for the cooperative treatment. In the 12 cooperative group sessions with inexperienced subjects there were 83 bidding periods in which all groups' cumulative balances were positive and 277 periods in which at least one group's balance was negative. In contrast, in the eight non-cooperative group sessions with inexperienced subjects there were 216 bidding periods in which all groups' cumulative balances were positive and 24 periods in which at least one group's balance was negative. In the 12 cooperative group sessions with once-experienced subjects there were 156 bidding periods in which all groups' cumulative balances were positive and 204 periods in which at least one group's balance was negative. In contrast, in the eight non-cooperative group sessions with once-experienced subjects there were 240 bidding periods in which all groups' cumulative balances were positive and no period in which any group's balance was negative.

4.2. Bidding Behavior by Nominal Groups

Table 2 reports summary comparisons of bidding behavior in all periods in which no bidding groups were bankrupt. The first column of Table 2 reports the experience of the groups (inexperienced = 0, once-experienced = 1). The second column shows the experimental treatment: cooperative or non-cooperative nominal bidding groups. The third column reports the average difference between individual subjects' signals and their bids. This is a measure of the extent to which individual subjects avoid the winner's curse by discounting their signals. The fourth column reports the standard deviation of the difference between individual subjects' signals and their bids. This is a measure of heterogeneity of individual subjects' discounting behavior. First consider the results for inexperienced subjects. The mean is higher and the standard deviation is lower in the non-cooperative treatment than in the cooperative treatment. Thus inexperienced subjects in the non-cooperative treatment are more effective in avoiding the winner's curse, and they are less heterogeneous in their discounting behavior than subjects in the cooperative treatment. Now compare the top two rows on Table 2 with the bottom two rows. Note that more experience leads bidders in both treatments to discount their signals by larger amounts. Finally, compare the bottom row of Table 2 with all other rows. Note that once-experienced subjects in the non-cooperative treatment discount their signals by the largest amount

Table 2. Bidders' Signal Discounts

Experience	*Group Type*	*Average Discount*	*Std. Dev. Discount*
0	Coop.	658	2240
0	Non-coop.	999	1890
1	Coop.	877	2544
1	Non-coop.	1270	1317

Table 3. Random Effects Regressions with Data for Nominal Groups

(standard errors)

Experience	*Group Type*	*Min. Rnl. Disc.*[a]	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\gamma}$	R^2	*Nobs.*[b]
0	Coop.	–900	156.48* (320.64)	0.9704[#] (0.014)	0.1502 (0.102)	0.92	264
0	Non-Coop.	–900	–265.17* (206.53)	0.9628[#] (0.009)	0.1416 (0.071)	0.94	630
1	Coop.	–900	–623.31* (78.0)	0.9901 (0.004)	0.171 (0.031)	0.99	585
1	Non-Coop.	–900	–949.31 (150.23)	0.9864 (0.007)	0.046 (0.048)	0.95	720

a. Min. Rnl. Disc. = minimum rational discount.
b. Nobs. = number of observations.
* Significantly greater than the minimum rational discount by a one-tailed 5% *t*-test.
[#] Significantly different than the theoretical value by a two-tailed 5% *t*-test.

and they are the least heterogeneous in this discounting behavior. Thus the "rationalizing" effect on bidding behavior of the free-riding incentive has a homogenizing effect on individual subjects' bidding behavior.

Table 2 includes all bids made by subjects in market periods with no bankrupt groups. We now turn our attention to analysis of market prices (winning bids). Table 3 reports results from random effects regressions with estimating equations of the form,

$$b_{jt} = \alpha + \beta m_{jt} + \gamma r_{jt} + \mu_j + \varepsilon_{jt}, \tag{6}$$

where b_{jt} is the bid by group j in period t, m_{jt} is group j's signal sample midrange in period t, and r_{jt} is group j's signal sample range in period t. The estimated coefficients are compared to the coefficients in the zero-expected-profit equation (5) to test for deviations from economic rationality. The estimation uses winning bids (market prices) and the associated right-hand variables.

The first and second columns of Table 3 report the experience level of the groups and the experimental treatment. The third column reports the "minimum rational discount," which is the intercept in the zero-expected-profit bid equation. The fourth, fifth and sixth columns report the estimated parameters and their standard errors (in parentheses). The seventh column reports the R^2's.

The last two rows in Table 3 report the random effects regression results for once-experienced subjects. Comparison of the estimated intercepts with the minimum rational discount and the estimated coefficients on slopes with a slope of 1 provides a measure of the departure from rational bidding by the high bidders in an experiment. The intercept for the cooperative-group treatment is −623, which is significantly greater than the minimum rational discount of −900 by a one-tailed t-test at the 5% confidence level, and the slope is 0.990, which is not significantly different from 1.000 by a two-tailed t-test at the 5% confidence level. Therefore, the winning bidders in the cooperative group treatment deviated significantly from minimally-rational bidding, in the direction of bidding to high; that is, cooperative bidding groups fell prey to the winner's curse. In contrast, the intercept for the non-cooperative group treatment is −949, which is obviously not greater than the minimum rational discount of −900, and the slope is 0.986, which is not significantly different from 1.000. Therefore, the winning bidders in the non-cooperative group treatment did not differ significantly from minimally-rational bidding; rather than falling prey to the winner's curse, the non-cooperative group bidders had positive expected profits. The incentive to free ride within non-cooperative groups tends to offset the winner's curse and promote rational bidding.

4.3. Comparison of Nominal and Natural Groups' Bidding Behavior

Table 4 reproduces data from two of the natural group treatments reported in Cox and Hayne (2002). Like the cooperative nominal group, the natural group treatment uses an equal profit-sharing rule. Unlike both cooperative and non-cooperative nominal groups, the natural group treatment involves face-to-face, within-group discussion and endogenously-determined rather than imposed decision rules.

Comparison of data from the two experiments can only be suggestive because the natural group experiments used five-member groups and the nominal group experiments used three-member groups. The first column of Table 4 reports the treatment parameters, Group size (5), Signal sample size (1 or 5) and Market size (3). Comparison of the intercept estimates in the two rows shows part of the support for the conclusion that more information (3 signals or common value estimates rather than 1) leads to less rational bidding by natural groups because the intercept estimate for the treatment with 5 signals (5, 5, 3) is significantly larger than the zero-profit

Table 4. Random Effects Regressions with Data for Natural Groups

(standard errors)

G, S, N[a]	*Min. Rnl. Disc.*[b]	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\gamma}$	R^2
5, 5, 3	–900	–527* (145)	0.994 (0.006)	0.154 (0.051)	0.998
5, 1, 3	–900	–708 (228)	0.984 (0.013)	–[c]	0.988

a. G, S, N = Group size, Signal sample size, Number of bidders.
b. Min. Rnl. Disc. = minimum rational discount
c. There is no estimated parameter for signal sample range here because the range is always zero by design in this treatment.
* Significantly greater than the minimum rational discount by a one-tailed 5% *t*-test.

intercept of –900 and the intercept estimate for the treatment with 1 signal (5, 1, 3) is not. Comparison of intercept estimates for natural (Table 4) and nominal (Table 3) groups leads to the following conclusions. Non-cooperative nominal group bidders are the only type that escapes the winner's curse and has bidding behavior with positive expected profits: –949 < –900. Cooperative nominal groups with 3 signals are less subject to the winner's curse than natural groups with 5 signals (–623 < –527) but more subject to the curse than natural groups with a single signal (–623 > –708).

5. CONCLUDING REMARKS

Data from our research on group bidding behavior supports some striking conclusions. The experiment reported in Cox and Hayne (2002) comparing bidding behavior of natural, face-to-face groups with bidding behavior by individuals reveals a "curse of information" that compounds the winner's curse. The bidding behavior of both individuals and natural groups deteriorates when they are given more information (a larger signal sample size) but bidding by groups deteriorates much more dramatically. Most strikingly, natural group bidders with more information (5 signals) are significantly less rational bidders than individuals with less information (1 signal).

Data from the nominal-group experiment reveal a rare instance in which an incentive to free ride leads to more, rather than less rational economic outcomes. The non-cooperative nominal group treatment, with the unequal profit-sharing rule providing a free-riding incentive, produced bidding behavior that was more rational than that observed with the cooperative nominal group treatment with no incentive to free riding.

ACKNOWLEDGMENT

Department of Economics, University of Arizona (Cox) and Department of Computer Information Systems, Colorado State University (Hayne). The authors are grateful for financial support from the Decision Risk and Management Science Program of the National Science Foundation (grant numbers SES-9709423 and SES-9818561) and to Rachel Croson and an anonymous referee for helpful comments on an earlier draft.

NOTES

[1] A deviation from this profit-sharing rule was required to handle some losses. If an individual member of a non-cooperative group attained a negative cumulative balance then other members of the group had to cover the loss. This was necessary to preclude a money pump that could result from limited liability of an individual subject. This cumulative loss-sharing rule was explained in the subject instructions.

[2] The experimental/U.S. dollar exchange rate was held constant across treatments.

[3] As shown by the subject instructions in the appendix, individual subjects were permitted to abstain rather than enter a bid. A few abstentions did occur, most by a single individual in a bankrupt group in the cooperative treatment.

REFERENCES

Capen, E., Clapp, R. & Campbell, W. (1971). "Competitive Bidding in High-Risk Situations," *Journal of Petroleum Technology*, 641–53.

Cox, J. & Hayne, S. (2002), "Barking Up the Right Tree: Are Small Groups Rational Agents?", The Behavioral Economics Conference, Great Barrington, MA, July 19–21.

Dyer, D. & Kagel, J. H. (1996). "Bidding in Common Value Auctions: How the Commercial Construction Industry Corrects for the Winner's Curse," *Management Science*, 42(10):1463–1475.

Hoffman, E., Marsden, J. & Saidi, R. (1991). "Are Joint Bidding and Competitive Common Value Auctions Markets Compatible – Some Evidence from Offshore Oil Auctions," *Journal of Environmental Economics and Management*, 20: 99–112.

Kagel, J. & Levin, D. (1986). "The Winner's Curse and Public Information in Common Value Auctions,"*American Economic Review*, 76(5): 894–920.

Kagel, J., Levin, D., Battalio, R. & Meyer, D. (1989). "First-Price Common Value Auctions: Bidder Behavior and the 'Winner's Curse'," *Economic Inquiry*, 27: 241–248.

APPENDIX. SUBJECT INSTRUCTIONS

A.1. Instructions for the Cooperative Nominal Group Treatment

Internet Auctions

INSTRUCTIONS

If you follow these instructions carefully, and make good decisions, you may earn a CONSIDERABLE AMOUNT OF MONEY. The amount of money you earn will be PAID TO YOU IN CASH at the end of the second day's experiment.

1. In this experiment we will create an auction market in which you will act as a member of a group bidding for a fictitious item in a sequence of many bidding periods. There will be several groups bidding on the item. A single unit of the

item will be auctioned off in each trading period. There will be several practice periods *without* money payoff followed by many "real" periods <u>with</u> money payoff.

2. Your task is to work with the other members of your groups and submit a group bid for the item. This will place your group in competition with other bidding groups. The precise value of the item at the time your group makes its bid will be unknown to you. Instead, each of you will receive a "signal" that provides an unbiased estimate of the item's value.

   Each individual in your group can submit a number that they think the group should bid bid (or an individual can abstain from bidding). The auction server computer will then **average** the numbers submitted by you and the other members of your group and submit that **average** as your group's bid in the auction. Abstentions are not included in this average.
3. When you bid in the auction, you will bid using *experimental* dollars. These experimental dollars can be redeemed at the end of the second day's experiment at the exchange rate shown on the computer. For example, if your group earned 4008 experimental dollars and the exchange rate was 80 experimental dollars per 1 U.S. dollar, then your group would earn $50.10 in **real U.S. dollars**.
4. The group with the highest bid in an auction period will be paid the value of the auctioned item and have to pay the amount of its bid. Thus, the group with the highest bid will receive a profit or loss equal to the difference between the value of the item and the amount that they bid:

GROUP PROFIT OR LOSS = VALUE OF ITEM − HIGHEST BID

If your group does not make the high bid on the item, your group will earn zero profit. In this case you neither gain nor lose from bidding on the item.

The group profit or loss is different from your individual profit or loss. Your individual profit or loss will be calculated by dividing the group profit or loss by the number of group members.

INDIVIDUAL PROFIT = GROUP PROFIT/GROUP SIZE

For example, if your group bid 12,885 experimental dollars (remember, this is the average of all the individual bids) for the object, it was higher than the other groups' bids and the value of the object was revealed to be 13,425 experimental dollars, your group profit would be 540 experimental dollars. If your group size was 3, then your individual profit would be 180 experimental dollars. You can see that if you work well in your group, you may earn a significant amount of money.

5. You will be given a starting capital credit balance of 1000 experimental dollars. Any profit you earn will be added to this amount and any losses will be subtracted from this amount. The net balance of these transactions will be calculated and paid to you in **CASH** at the end of the second day's experiment. The starting capital credit balance, and whatever subsequent profits you earn, permit you to suffer losses in one auction that could be recouped in part or total in later auctions.

Your group is permitted to bid in excess of your own capital credit balance in any given period.

6. During each trading period, your group will be bidding in a market with several other groups and after all the bids have been submitted, the winning bid will be announced.
7. The value of the item will be chosen randomly each auction period and will always lie between 2,500 and 22,500 experimental dollars, inclusively. For each auction, any value within this interval has an equally likely chance of being drawn. The value of the item can never be less than 2,500 or more than 22,500 experimental dollars. The values are determined randomly and are independent from auction to auction. As such, a high value in one auction tells you nothing about the likely value in the next auction, i.e. whether it will be high or low.
8. Private Information Signals: Although you do not know the precise value of the item in any particular auction, you will receive information which will narrow down the range of possible values. This will consist of a private information signal which is selected randomly from an interval whose lower bound is the item value less a constant amount, and whose upper bound is the item value plus the same constant. Any value within this interval has an equally likely chance of being drawn and being assigned to you as your private information signal. The value of this constant will be announced prior to the experiment.

For example, suppose that the value of the auctioned item is 12,677 experimental dollars and that the constant is 1,800 experimental dollars. Each of you will receive a private information signal which will consist of a randomly drawn number that will be between 10,877 (12,677 – 1,800) and 14,477 (12,677 + 1,800) experimental dollars. Any number in this range has an equally likely chance of being drawn.

The data below shows an entire set of signals the computer might generate for a group of ten people. (Note these have been ordered from highest to lowest).

The item value is 12,677 and the constant is 1,800 experimental dollars, and the signals are:

14314
13730
13709
13331
12917
12435
12344
11971
11785
11385

You can see that some signal values were above the value of the auctioned item, and some were below the value of the item. Over a sufficiently long series of auctions, the differences between your private signals and the item values will average out to zero (or very close to it). But for any single auction your private information signal can be above or below the value of the item. That's the nature of the random drawing process that is generating the signals. You will also note that the upper bound must always be greater than or equal to your signal value. Further, the lower bound must always be less than or equal to your signal value.

Finally, you may receive a signal value below 2,500 (or above 22,500). There is nothing strange about this, it merely indicates that the item value is close to 2,500 (or 22,500) and this closeness depends on the size of the constant.

9. Your signals are strictly private information.
10. Bids are rounded to the nearest experimental dollar and must be greater than 0. In case of ties for the high bid, a coin toss will determine the winner.
11. You are not to communicate with anyone while the experiment is in progress.

SUMMARY OF MAIN POINTS

1. A group's bid is the average of the bids submitted by individual members of the group. The group with the highest bid wins the auction and receives a profit or loss.
2. Your individual profit or loss will equal your group's profit or loss divided by group size.
3. Your cumulative profit will be paid to you, in **CASH**, at the end of the second day's experiment.
4. Your private information signal is drawn from the interval (item value – 1,800, item value + 1,800). The value of the item can be as much as 1,800 below your signal or 1,800 above your signal.
5. The value of the item will always be between 2,500 and 22,500.

ARE THERE ANY QUESTIONS?

Part 2. Instructions for the Non-cooperative Nominal Group Treatment

The instructions for the non-cooperative treatment were the same as for the cooperative treatment, except as explained here. Paragraph 4 in the INSTRUCTIONS (first part) of section A.1 was replaced by the following paragraphs 4 and 5. (Paragraphs 6–12 in section A.2 are the same as paragraphs 5–11 in section A.1.) The SUMMARY OF MAIN POINTS in section A.1 was replaced by the one below.

4. If your group does not make the highest bid on the item, each member of your group will receive zero profit or loss. The group with the highest bid in an auction period will be paid the value of the auctioned item and have to pay the amount of its bid.

The individual members of the group with the highest bid do **not** usually share the profit or loss equally. If your group has the highest bid, your individual profit or loss will be calculated by subtracting your bid from the value of the item and dividing that profit or loss by three, the number of people in your group:

$$\text{INDIVIDUAL PROFIT OR LOSS} = \frac{\text{ITEM VALUE} - \text{YOUR BID}}{3}$$

For example, suppose that your group has the highest bid and the value of the object turns out to be 13,425 experimental dollars. If your individual bid was 10,560 then your individual profit would be 955 experimental dollars ((13425–10560)/3). However, if your individual bid was 15,220 then your individual profit would be –598; a **loss** of your experimental dollars ((13425–15220)/3). There is an exception to the above way of calculating individual profits that occurs if any member of your group becomes bankrupt, that is if someone attains negative total payoff.

5. If you abstain from bidding in any period, and your group has the highest bid, then your profit or loss will equal one third of the difference between the item value and the average of the bids submitted by other members of your group.

SUMMARY OF MAIN POINTS

1. A group's bid is the average of the bids submitted by individual members of the group. The group with the highest bid wins the auction and receives a profit or loss.
2. If your group has the highest bid, your individual profit or loss will be equal to 1/3 of the difference between the item value and your bid, so long as no one in your group is bankrupt.
3. If the total payoff of someone in your group becomes negative then the other members of the group must cover that person's losses until such time as he/she attains positive total payoff.
4. Your total payoff will be paid to you, in **CASH**, at the end of the second day's experiment.
5. Your private information signal is drawn from the interval (item value – 1,800, item value + 1,800). The value of the item can be as much as 1,800 below your signal or 1,800 above your signal.
6. The value of the item will always be between 2,500 and 22,500.

ARE THERE ANY QUESTIONS?

Chapter 8

BONUS VERSUS PENALTY: DOES CONTRACT FRAME AFFECT EMPLOYEE EFFORT?

R. Lynn Hannan
Georgia State University

Vicky B. Hoffman
University of Pittsburgh

Donald V. Moser
University of Pittsburgh

Abstract

We conducted an experiment in which participants acted as employees under either a bonus contract or an economically equivalent penalty contract. We measured participants' contract preference, their degree of expected disappointment about having to pay the penalty or not receiving the bonus, their perceived fairness of their contract, and their effort level. Consistent with Luft (1994), we find that employees generally preferred the bonus contract to the penalty contract. We extend Luft's work by demonstrating that loss aversion caused employees to expend more effort under the penalty contract than under the economically equivalent bonus contract. That is, employees were more averse to having to pay the penalty than they were to not receiving the bonus, and consequently they chose a higher level of effort under the penalty contract to avoid paying the penalty. However, we also find evidence of reciprocity in that employees who considered their contract to be fairer chose a higher level of effort. Because our participants generally considered the bonus contract fairer than the penalty contract, reciprocity predicts higher effort under the bonus contract, a result opposite to our finding. Our overall result that employee effort was greater under the penalty contract is explained by the fact that, while higher perceived fairness did increase effort, this effect was dominated by the more powerful opposing effect of loss aversion. We discuss the implications of these results for explaining why in practice most actual contracts are bonus contracts rather than penalty contracts.

A. Rapoport and R. Zwick (eds.), Experimental Business Research, Vol. II, 151–169.

1. INTRODUCTION

Conventional economic analysis of incentive contracts in managerial accounting settings (e.g., Demski and Feltham 1978; Holmstrom 1979, 1982; Holmstrom and Milgrom 1991; Feltham and Xie 1994) assumes that utility increases in wealth and decreases in effort. Moreover, cognitive factors such as framing and loss aversion and social norms such as fairness and reciprocity are assumed to be relatively unimportant, and as such, are ignored in the analysis. Given that conventional economic analysis predicts that the way economically equivalent contracts are framed should not matter, an unsolved puzzle is why in practice most incentive contracts are framed as bonus contracts rather than penalty contracts (Baker, Jensen & Murphy 1988; Milgrom & Roberts 1992; Young & Lewis 1995). By "economically equivalent" we mean that the monetary incentives are identical under both the bonus and penalty versions of a contract. For example, a bonus contract that pays a salary of \$20 and a bonus of \$10 if a target outcome is achieved is equivalent to a penalty contract that pays a salary of \$30 and a penalty of \$10 if the target outcome is not achieved. These two contracts are economically equivalent because under both contracts the employee receives \$30 if the target outcome is achieved and \$20 if the target outcome is not achieved.

Luft (1994) offered a potential explanation for why in practice most contracts are framed in bonus terms by demonstrating that employees are more likely to choose economically equivalent contracts that are framed as bonus contracts rather than penalty contracts. She attributed employees' preferences for bonus contracts to loss aversion (Kahneman & Tversky 1979), arguing that employees experience greater disutility from the perceived loss associated with paying a penalty than from the perceived forgone gain associated with not receiving an equivalent bonus. If loss aversion causes employees to demand higher payments from firms to accept penalty contracts, then firms maximize profits by offering bonus contracts.[1] However, this explanation for why most actual incentive contracts are framed as bonus contracts assumes that firms can offer bonus contracts at no greater cost than economically equivalent penalty contracts. A possible cost of offering bonus contracts that is not reflected in either the standard economic analysis or in Luft's explanation is lower employee effort. Bonus contracts could yield lower effort than economically equivalent penalty contracts if employees facing penalty contracts expend greater effort to avoid the perceived loss associated with the potential penalty. That is, the same theoretical construct of loss aversion that predicts employees will choose bonus contracts over economically equivalent penalty contracts, also predicts that employees will expend more effort under penalty contracts. Moreover, the desire to avoid the pressure to expend more effort under penalty contracts may help explain why employees prefer bonus contracts to economically equivalent penalty contracts.

If employees generally expend more effort under penalty contracts than under economically equivalent bonus contracts, an important implication is that the preference for bonus contracts documented by Luft no longer necessarily provides an explanation for why most actual contracts are framed in bonus terms. That is, if employees expend more effort under penalty contracts, employers would then need to trade off the higher level of employee effort associated with penalty contracts (a benefit) against the higher payments required to induce employees to accept penalty contracts (a cost) in order to determine whether a bonus contract or a penalty contract will maximize firm profit. Luft's study could not address this issue directly because it was not designed to examine whether framing an incentive contract as a bonus contract versus as a penalty contract differentially motivated employee effort. Specifically, Luft's incentive contract was designed to improve outcomes by separating agent types rather than by motivating effort. Her participants performed a recognition memory task that was relatively insensitive to effort (i.e., performance depended mostly on prior knowledge rather than effort). Consequently, performance did not differ across her bonus and penalty conditions.

Interestingly, responses to Luft's post-experimental questionnaire suggest another reason in addition to loss aversion for why employees prefer bonus contracts to penalty contracts. That is, virtually all participants indicated that they thought that "most employees" would feel that a bonus scheme was fairer than an economically equivalent penalty scheme. Thus, employees' preferences for bonus contracts could be due, in part, to their preference for working under a fairer contract. However, in contrast to loss aversion, the theory of reciprocity (Rabin 1993) predicts that employee effort will be higher under bonus contracts. That is, this alternative theory predicts that if employees view bonus contacts as fairer, they will reciprocate by expending more effort under bonus contracts than under penalty contracts. Of course, this prediction is opposite to the prediction that employees will expend more effort under penalty contracts because of loss aversion

To recap, both loss aversion and perceived fairness predict that employees will prefer bonus contracts to economically equivalent penalty contracts. However, loss aversion and reciprocity make opposite predictions regarding employee effort. With two forces pushing effort in opposite directions, it is not clear whether employee effort will be higher, lower, or about the same under bonus versus economically equivalent penalty contracts. Therefore, before we can explain why most actual incentive contracts are framed as bonus contracts rather than penalty contracts, we need a better understanding of the factors underlying employees' preferences for bonus contracts and whether and how these factors affect employees' effort. The purpose of this study is to help provide such an understanding.

2. HYPOTHESES AND RESEARCH QUESTIONS

Our first hypothesis deals with employees' preferences. As discussed above, Luft (1994) found that employees prefer bonus contracts to penalty contracts. Given Luft's compelling theoretical arguments and her strong experimental results, we expect to replicate her finding. Thus, like Luft, we predict that employees will prefer bonus contracts.

H1: Employees prefer bonus contracts to economically equivalent penalty contracts.

Before we address expected differences in effort between bonus and penalty contracts, we first hypothesize general effects in our second and third hypotheses that we expect to hold for both bonus and penalty contracts. Our second hypothesis addresses the effect on effort of employees' expected disappointment about having to pay a penalty or not receiving a bonus. We do not distinguish between bonus contracts and penalty contracts because disappointment is expected to affect effort regardless of whether the contract is framed as a bonus or as a penalty. Specifically, we predict that employees who expect to feel more disappointed about the prospect of receiving lower compensation (either by having to pay a penalty or by not receiving a bonus) will expend more effort to avoid that outcome than employees who expect to feel less disappointed about receiving the lower final payment. This prediction is consistent with conventional economic theory, which assumes that employees with greater incremental utilities for a higher monetary outcome (i.e., receiving the higher final payment without having to pay a penalty or forgo a bonus) will expend more effort to ensure that they receive that outcome. Thus, it follows that employees with a greater incremental utility for receiving a higher monetary outcome will experience a greater reduction in utility from not receiving that outcome. In our study, "expected disappointment" about not receiving the bonus or having to pay the penalty corresponds to this reduction in utility from not receiving the higher final payment.

H2: Greater expected disappointment will result in higher employee effort.

Our third hypothesis relates to the effect of perceived fairness on effort. Many studies in psychology (e.g., Goranson and Berkowitz 1966; Greenberg and Frisch 1972; Greenberg 1978) and experimental economics (e.g., Kahneman, Knetsch and Thaler 1986; Fehr, Kirchsteiger and Riedl 1993; Fehr, Gächter and Kirchsteiger 1997; Fehr, Kirchler, Weichbold and Gächter 1998; Charness and Rabin 2002; Hannan, Kagel and Moser 2002) have shown that individuals who feel they are treated fairly by another party will reciprocate by treating that party kindly in return. This theory of "reciprocity" underlies our third hypothesis, which predicts that employees who perceive their contract to be fairer will choose a higher level of effort than those who perceive their contract to be less fair. As was the case for H2, this is a general hypothesis that does not distinguish between bonus contracts and penalty contracts. That is, higher perceived fairness is predicted to yield higher employee effort in both bonus contracts and penalty contracts.

H3: Employees who perceive their contracts to be fairer will expend higher effort.

As explained above, the general effects of expected disappointment (H2) and perceived fairness (H3) on effort are predicted to operate in the same manner within both bonus contracts and penalty contracts. However, as discussed below, the levels of both expected disappointment and perceived fairness are likely to differ *across* bonus and penalty contracts.

With respect to disappointment, the theoretical construct of loss aversion predicts that expected disappointment would be greater under penalty contracts than under economically equivalent bonus contracts. Loss aversion describes the well-documented finding that individuals are more averse to suffering a loss than they are to forgoing the same amount of gain (Kahneman and Tversky 1979). If employees facing penalty contracts frame the prospect of having to pay the penalty as a loss, they will expect to be very disappointed about having to pay the penalty. In contrast, if employees facing bonus contracts frame the prospect of not receiving an economically equivalent bonus as a forgone gain, they will expect to be less disappointed about not receiving the bonus. These asymmetric framing effects across contract type lead to our fourth hypothesis.

H4: Employees facing a penalty contract will expect to be more disappointed about having to pay a penalty than employees facing a bonus contract will expect to be about not receiving an economically equivalent bonus.

If greater disappointment results in more effort (H2), and disappointment is greater under penalty contracts than under bonus contracts (H4), then it follows that employee effort should be greater under penalty contracts than under bonus contracts. However, as explained below, the fact that reciprocity predicts an opposing effect on effort prevents us from making such a simple directional prediction regarding the effect of contract frame (bonus or penalty) on employee effort.

With respect to perceived fairness, virtually all of Luft's (1994) participants indicated in her post-experimental questionnaire that they thought that "most employees" would perceive a bonus contract to be fairer than an economically equivalent penalty contract. Such perceptions could be due to a construct that Luft refers to as "penalty aversion." If employees are averse to penalty contracts because they view penalty contracts as punitive or negative, they are likely to perceive penalty contracts as unfair. In contrast, if bonus contracts are viewed more positively because employees frame them as offering a potential reward, they are likely to be perceived as fairer than economically equivalent penalty contracts. These expected differences across contract types lead to our fifth hypothesis.

H5: Employees perceive bonus contracts as fairer than economically equivalent penalty contracts.

Hypotheses H2-H5 are depicted in Figure 1, where it can be seen that if employees consider bonus contracts to be fairer than penalty contracts (H5) and also engage in

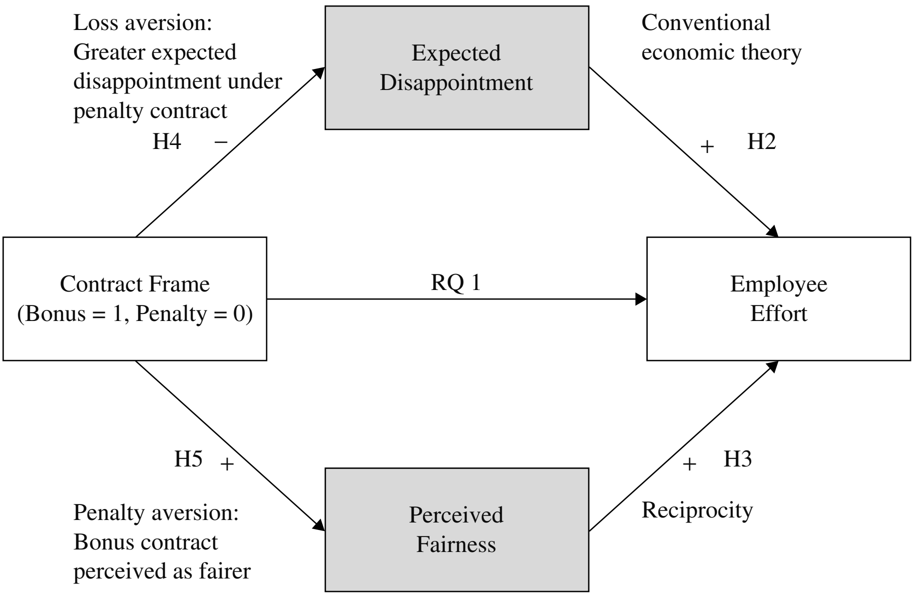

Figure 1.

reciprocity (H3), then it follows that employee effort should be higher under bonus contracts than under economically equivalent penalty contracts (bottom path in Figure 1). Of course, this prediction regarding employee effort is opposite to the prediction described earlier that effort will be higher under penalty contracts (top path in Figure 1) as a result of the combined effect of loss aversion (H4) and expected disappointment (H2). Because these potential effects work in opposite directions, we are unable to predict the net effect on effort of framing contracts in bonus versus penalty terms. Therefore, we do not make a directional hypothesis regarding the effect of contract frame on effort, but rather address this issue as our first research question (RQ1 in Figure 1).

RQ1: Does employee effort differ under economically equivalent contracts framed in bonus versus penalty terms, and if so, which type of contract results in higher effort?

We expand upon RQ1, by investigating a second research question, RQ2 (not directly identified in Figure 1), which involves isolating and measuring the potentially offsetting effects of loss aversion and perceived fairness on effort. Specifically, RQ2 addresses whether expected disappointment, perceived fairness, or both expected disappointment and perceived fairness mediate the effect of contract frame (bonus versus penalty) on employee effort. As explained earlier, if H4 and H2 (top path in Figure 1) are supported, then contract frame is likely to affect effort by way of expected disappointment. However, if H5 and H3 (bottom path in Figure 1) are also supported, then it is likely that contract frame also affects effort by way of perceived fairness. Because we cannot predict in advance whether expected disappointment, perceived fairness, or both will mediate the effect of contract frame on effort, we address these issues in our second research question.

RQ2: Does expected disappointment or perceived fairness mediate any effect of contract frame on employee effort (examined in RQ1)?

3. EXPERIMENT

3.1. Overview

We conducted an experiment designed to address the hypotheses and research questions described above. Participants were assigned to either a bonus contract or an economically equivalent penalty contract (described later). Their task was to choose their effort level. In addition, they responded to questions designed to measure their degree of expected disappointment about not receiving the bonus or having to pay the penalty and their perceived fairness of the contract they faced when making their effort choices. After making their effort choices and responding to the expected disappointment and fairness questions, participants were shown the contract that participants in the other experimental condition faced (i.e., bonus contract participants were shown the penalty contract, and vice versa) and asked to indicate which contract they preferred.

3.2. Participants

Sixty-eight M.B.A. students participated in the experiment. Sixty-two percent of the participants had at least five years of professional business experience, with the remainder having between zero and five years of experience. Forty-seven percent of participants had worked under a bonus incentive contract. No participants had worked under a penalty incentive contract. Both professional business experience and incentive contract experience were distributed approximately equally across the experimental (bonus and penalty) conditions.

3.3. Design

Our experimental design included a manipulated between-subjects independent variable, Contract Frame, with two levels (Bonus and Penalty). The design also included two measured variables (Expected Disappointment and Perceived Fairness) that were obtained from participants' responses to questions in the experimental instrument. Finally, our design included two dependent variables: participants' effort level choices and their expressed contract preference. As explained in the results section of the paper, the specific combination of independent and dependent variables used for any particular analysis depended on the hypothesis or research question being addressed.

3.4. Procedures

The experiment was conducted in two back-to-back administrations, one for each experimental condition. Each administration took approximately 30 minutes. Participants were randomly assigned to either the Bonus or economically equivalent Penalty condition. The bonus contract paid a salary of \$20 plus a bonus of \$10 if the target (high) outcome was achieved. The economically equivalent penalty contract paid a salary of \$30 with a \$10 penalty if the target (high) outcome was not achieved. These contracts are economically equivalent because under both contracts the employees will receive \$30 if the outcome is high and \$20 if the outcome is low. Participants assigned to either condition were unaware that the alternative condition existed until after they made their effort choices and responded to the expected disappointment and perceived fairness questions.

Participants in both conditions assumed the role of an employee of Buckley Company. They received their base pay in cash (\$20 in the bonus condition, \$30 in the penalty condition) at the start of the experiment, and were told that their final payment at the end of the experiment (i.e., the cash they retained or the additional cash they received) would depend on the terms of their contract and the effort level they chose.[2] The description of Buckley Company indicated that the company's goal was to maximize shareholder value. Company management had instituted a new compensation system designed to provide an incentive for employees to work hard to achieve a high outcome so that the company could meet its aggressive performance goals. The more effort an employee expended, the more likely it was that s/he would achieve a high outcome.

Consistent with previous studies (e.g., Frederickson 1992; Fehr, Kirchsteiger and Riedl 1993; Fehr, Gächter and Kirchsteiger 1997; Hannan, Kagel and Moser 2002) disutility for effort was operationalized as a monetary cost to participants that increased with the level of effort chosen.[3] Specifically, participants chose an effort level from 1 to 13, with the cost of effort increasing correspondingly in \$.50 increments from \$.50 (1) to \$6.50 (13). The probability of achieving a high outcome also increased with the level of effort, rising in 5% increments from 30% (1) to 90% (13). The cost of effort and probabilities of achieving a high outcome were set such that the participants' expected net payoff was identical (\$22.50) across all 13 possible effort level choices.[4]

Participants were provided a table that showed the cost of effort and the probability of achieving (and not achieving) the high outcome for each of the 13 possible effort level choices (see Table 1). After reading a description of their employment contract and reviewing this table, each participant chose his or her effort level. Immediately after making their effort level choices, participants responded to the fairness and expected disappointment questions. Participants rated the fairness of the employment contract they faced in the experiment on a 13-point scale with endpoints labeled "not fair at all" (1), and "extremely fair" (13), and the midpoint labeled "moderately fair" (7). Participants rated how disappointed they would be if the outcome were low and therefore they did not receive the bonus (had to pay the penalty)

Table 1. Cost of Effort Tables for Penalty and Bonus Contract Frames

Penalty Contract Frame

Your effort level	*Cost of Effort*	*Probability of Achieving a High Outcome and Avoiding the $10 Penalty*	*Probability of not Achieving a High Outcome and Paying the $10 Penalty*
1	$.50	30%	70%
2	$1.00	35%	65%
3	$1.50	40%	60%
4	$2.00	45%	55%
5	$2.50	50%	50%
6	$3.00	55%	45%
7	$3.50	60%	40%
8	$4.00	65%	35%
9	$4.50	70%	30%
10	$5.00	75%	25%
11	$5.50	80%	20%
12	$6.00	85%	15%
13	$6.50	90%	10%

on a 13-point scale with endpoints labeled "not at all disappointed" (1), and "extremely disappointed" (13), and the midpoint labeled "moderately disappointed"(7).

After responding to the fairness and expected disappointment questions, participants were provided with a description of the employment contract used in the other condition (i.e., the bonus condition participants now saw the penalty contract, and vice versa) and the related effort-choice table. After considering this information, participants indicated whether they preferred the original contract they faced in the experiment, the alternative contract they were considering now, or had no preference between the two. Responses to this question were used to test whether most of our participants preferred the bonus contract to the economically equivalent penalty contract, irrespective of whether they were assigned to the bonus or penalty condition. The experimental instrument concluded with several demographic questions

Table 1. (cont'd)

Bonus Contract Frame

Your effort level	*Cost of Effort*	*Probability of Achieving a High Outcome and Receiving the $10 Bonus*	*Probability of not Achieving a High Outcome and Not Receiving the $10 Bonus*
1	$ .50	30%	70%
2	$1.00	35%	65%
3	$1.50	40%	60%
4	$2.00	45%	55%
5	$2.50	50%	50%
6	$3.00	55%	45%
7	$3.50	60%	40%
8	$4.00	65%	35%
9	$4.50	70%	30%
10	$5.00	75%	25%
11	$5.50	80%	20%
12	$6.00	85%	15%
13	$6.50	90%	10%

regarding participants' professional work experience and their experience with incentive contracts.

After all participants had submitted their experimental materials, the actual outcome (high or low) was determined for each effort level (1 through 13) as follows: A participant volunteer drew one chip from each of 13 bags (one for each effort level). Each bag contained red and blue chips in proportion to the outcome probability distribution corresponding to that effort level. For example, because effort level 1 had a 30% probability of a high outcome and 70% probability of a low outcome, the bag for effort level 1 contained 3 red chips (high outcome) and 7 blue chips (low outcome).[5] After outcomes had been determined for each effort level, participants' final payments were calculated and they were paid in cash privately. Bonus condition participants who did not receive a bonus (outcome was low) were required to return a portion of their \$20 base pay equal to the cost of their chosen effort level. Bonus condition participants who received a bonus (outcome was high) were paid an additional sum equal to the \$10 bonus minus the cost of their chosen effort level. Penalty condition participants who had to pay the penalty (outcome was low) were required to return a portion of their \$30 base pay equal to the cost of their effort plus the \$10 penalty. Penalty condition participants who did not have to pay the penalty (outcome was high) were required to return a portion of their \$30 base pay equal to the cost of their chosen effort level.

4. RESULTS

4.1. Tests of Hypotheses 1–5

H1 predicts that employees prefer bonus contracts to economically equivalent penalty contracts. To test this hypothesis we examined participants' preference responses after they considered both the original contract they faced in the experiment and the contract used in the other experimental condition (i.e., after bonus participants were provided with the penalty contract, and vice versa). Overall, 65% of participants preferred the bonus contract, 19% preferred the penalty contract, and 16% were indifferent between the two. Although conventional economic theory predicts that all participants would be indifferent, 84% of participants expressed a preference, and a significantly greater proportion (binomial test, $p < .001$) of these preferred the bonus contract (65%) to the penalty contract (19%). Results were not significantly different across experimental conditions (chi square = 1.47, $p = .48$), with 60% (70%) preferring the bonus contract, 23% (15%) preferring the penalty contract, and 17% (15%) expressing indifference between the contracts in the bonus and penalty conditions, respectively. Further, of those participants expressing a preference, the proportion preferring the bonus contract in each experimental condition was significantly greater than the proportion preferring the penalty contract in that condition (binomial tests, $ps < .03$). These results are consistent with H1, and as such, replicate Luft's finding that employees generally prefer bonus contracts to penalty contracts.

H2 predicts that greater expected disappointment about having to pay the penalty or not receiving the bonus will result in greater employee effort in both the bonus and penalty conditions. To test this hypothesis, we first regressed participants' effort choices on their expected disappointment responses for the pooled data set. Results show a strong positive association ($t = 5.61$, $p < .001$), indicating that, consistent with H2, effort increased significantly as expected disappointment increased. Separate regressions for the Bonus ($t = 3.67$, $p < .002$) and Penalty ($t = 3.49$, $p < .002$) conditions yielded similar results.

H3 predicts that higher employee fairness ratings of their contracts will result in higher employee effort in both the bonus and penalty conditions. To test this hypothesis we first regressed participants' effort choices on their perceived fairness ratings for the pooled data set. Results showed a strong positive association (t = 2.79, $p < .008$), indicating that, consistent with H3, effort increased significantly as

Table 2. Mean (standard deviation) of Expected Disappointment, Perceived Fairness and Effort Measures

Variable	*Contract Frame*		*t-statistic (Bonus = Penalty)*	*p-value (two tailed)*
	Bonus	*Penalty*		
Expected Disappointment	6.77 (3.26)	9.36 (3.51)	3.16	.002
Perceived Fairness	7.40 (3.94)	5.30 (3.17)	2.41	.019
Effort	7.40 (4.58)	9.58 (3.31)	2.23	.029
n	35	33		

perceived fairness increased. Separate regressions for the Bonus ($t = 3.13$, $p < .005$) and Penalty ($t = 2.08$, $p < .047$) conditions yielded similar results.

H4 predicts that employees facing a penalty contract will expect to be more disappointed about having to pay the penalty than employees facing a bonus contract will be about not receiving the bonus. In other words, participants' expected disappointment will be asymmetric, reflecting loss aversion. To test this hypothesis we compared participants' ratings of the degree of disappointment they expected to feel if they did not receive the bonus in the Bonus condition to the degree of disappointment they expected to feel if they had to pay the penalty in the Penalty condition. The results, which are reported in Table 2, show that, consistent with H4, participants' degree of expected disappointment was significantly higher ($t = 3.16$, $p < .003$) in the Penalty condition (mean = 9.36) than in the Bonus condition (mean = 6.77). That is, despite the economic equivalence of the bonus and penalty contracts, participants indicated that they were significantly more averse to having to pay the penalty than they were to not receiving the bonus. These results reflect loss aversion because participants viewed paying the penalty as a bigger psychological loss than not receiving the bonus.

H5 predicts that employees will perceive bonus contracts to be fairer than penalty contracts. To test this hypothesis, we compared participants' ratings of how fair they considered the contract they faced in the Bonus versus Penalty conditions. As shown in Table 2, Bonus condition participants rated the bonus contract (mean = 7.40) as significantly fairer ($t = 2.41$, $p < .02$) than Penalty condition participants rated the penalty contract (mean = 5.30). These results are consistent with H5, as well as with Luft's post-experimental questionnaire results, which showed that virtually all of her participants thought that "most employees" would feel that a bonus contract was fairer than an economically equivalent penalty contract.

4.2. Research Questions

Our first research question (RQ1) asks whether framing economically equivalent contracts in bonus terms versus in penalty terms affects employee effort. As shown in Table 2, employee effort was significantly higher (t = 2.23, p = .029 two-tailed) in the Penalty condition (mean = 9.58) than in the Bonus condition (mean = 7.40). This result can potentially be explained by the loss aversion documented earlier in tests of H4, which indicated that Penalty condition participants expected to be more disappointed about having to pay the penalty than Bonus condition participants expected to be about not receiving the bonus. Combined with the finding that greater disappointment resulted in higher employee effort (H2), these results can explain why employee effort was greater in the penalty condition.[6]

The greater effort observed under the penalty contract runs contrary to a reciprocity effect which predicts that effort will be greater under the bonus contract. Nevertheless, reciprocity could still be operating if the effect were dominated by the more powerful opposing effect of loss aversion. Indeed, further analysis reported below for our second research question is consistent with this interpretation.

Our second research question (RQ2) asks whether Expected Disappointment and/or Perceived Fairness mediate the effect of Contract Frame on Effort documented in RQ1. To address this question, we conducted four regression analyses as follows:

(1) Effort = $\alpha + \beta_1$ Contract Frame + ε
(2) Effort = $\alpha + \beta_1$ Contract Frame + β_2 Expected Disappointment + ε
(3) Effort = $\alpha + \beta_1$ Contract Frame + β_2 Perceived Fairness + ε
(4) Effort = $\alpha + \beta_1$ Contract Frame + β_2 Expected Disappointment + β_3 Perceived Fairness + ε

where, Effort = participants' effort choices
Contract Frame = 1 for Bonus condition, 0 for Penalty condition
Expected Disappointment = participants' rating of the disappointment they expected to experience if they did not receive the bonus (Bonus condition) or had to pay the penalty (Penalty condition)
Perceived Fairness = participants' rating of the fairness of their contract

The results for these four regressions are reported in Table 3. We know from the analysis reported for RQ1 that, overall, Effort was higher in the Penalty condition (Contract Frame = 0) than in the Bonus condition (Contract Frame = 1). This is confirmed by the results of the first regression, which show that Contract Frame is negatively related to Effort (t = –2.23, p = .029). The second regression examines the extent to which the effect of Contract Frame on Effort is mediated by Expected Disappointment. The results indicate that, consistent with the results of H2 reported earlier, Expected Disappointment has a strong positive effect (t = 4.97, p < .001) on Effort. However, more importantly, including Expected Disappointment as an

Table 3. Effort Regressions

		Intercept	*Contract Frame*	*Expected Disappointment*	*Perceived Fairness*	*Adj. R^2*
Regression Model 1	Coefficient (standard error) *t*-statistic *p-value*	9.576 (0.699) 13.70 .000	−2.176 (0.974) −2.23 .029			.06
Regression Model 2	Coefficient (standard error) *t*-statistic *p-value*	3.734 (1.319) 2.83 .006	−0.559 (0.897) −.62 .536	0.624 (0.125) 4.97 .000		.31
Regression Model 3	Coefficient (standard error) *t*-statistic *p-value*	7.006 (0.919) 7.62 .000	−3.192 (0.924) −3.46 .001		0.485 (0.125) 3.87 .000	.22
Regression Model 4	Coefficient (standard error) *t*-statistic *p-value*	1.170 (1.268) .92 .360	−1.575 (0.808) −1.95 .056	0.624 (0.109) 5.72 .000	0.484 (0.103) 4.71 .000	.48

Number of observations = 68 for all models.

The full model (Model 4) is:

Effort = $\alpha + \beta_1$ Contract Frame + β_2 Expected Disappointment + β_3 Perceived Fairness + ε

where, Effort = participants' effort choices
Contract Frame = 1 for Bonus condition, 0 for Penalty condition
Expected Disappointment = participants' rating of the disappointment they expected to experience if they did not receive the bonus (Bonus condition) or had to pay the penalty (Penalty condition)
Perceived Fairness = participants' rating of the fairness of their contract

explanatory variable in the second regression causes the effect of Contract Frame on Effort to drop to nonsignificance ($t = -.62$, $p = .536$ in the second regression versus $t = -2.23$, $p = .029$ in the first regression). These results indicate that Expected Disappointment mediates the effect of Contract Frame on Effort. That is, the reason that employees chose greater effort in the Penalty condition than in the Bonus condition appears to be that, consistent with the loss aversion documented in H4, they were more averse to having to pay the penalty in the Penalty condition than they were to not getting the bonus in the Bonus Condition. Moreover, the adjusted R^2 increased substantially when Expected Disappointment was included in the second regression (R^2 increased from .06 in the first regression to .31 in the second regression), showing that not only does Expected Disappointment mediate the effect of Contract on

Effort, but it also has a strong, separate influence on Effort within each of the Bonus and Penalty conditions. These results are consistent with the top path in Figure 1.

The third regression shows that Perceived Fairness also mediates the effect of Contract Frame on Effort (bottom path of Figure 1), but in the opposite direction of Expected Disappointment. Consistent with the reciprocity documented earlier in H3, Perceived Fairness has a strong positive effect ($t = 3.87$, $p < .001$) on Effort, indicating that higher Perceived Fairness yields higher Effort. In addition, after statistically controlling for the effect of Perceived Fairness, the effect of Contract Frame on Effort was actually stronger ($t = -3.46$, $p < .002$) than it was in the first regression ($t = -2.23$, $p = .029$) when Perceived Fairness was omitted from the regression. That is, including Perceived Fairness in the regression does not weaken the effect of Contract Frame on Effort as is the case in most mediation analysis, but rather strengthens it. This is because, although both Expected Disappointment and Perceived Fairness mediate the effect of Contract Frame on Effort, they do so in opposite directions.[7] A final important observation regarding the third regression is the substantial increase in the adjusted R^2 from .06 in the first regression to .22 in the third regression, which indicates that in addition to mediating the effect of Contract Frame on Effort, Perceived Fairness had a separate effect on Effort within each of the Contract Frame conditions.

Although Perceived Fairness was higher in the Bonus condition than in the Penalty condition (H5) and it affected Effort as predicted in H3 (i.e., higher perceived fairness led to greater effort), this effect (depicted in the bottom path of Figure 1) was entirely dominated by the opposing effect depicted in the top path of Figure 1. Taken together, the results of the first three regressions suggest that, while reciprocity caused participants to choose more effort in the Bonus condition than in the Penalty condition (third regression), this effect was dominated by the stronger opposing effect of loss aversion (second regression), which caused participants to choose more effort under the penalty contract (first regression).

The fourth regression, which included both Perceived Fairness and Expected Disappointment as explanatory variables, increased the adjusted R^2 to .48, versus .22 when only Perceived Fairness was included in the third regression, and .31 when only Expected Disappointment was included in the second regression. As shown in Table 3, the effects of both Perceived Fairness and Expected Disappointment on Effort remain statistically significant in the expected directions in the fourth regression. These results confirm the interpretation offered above, which was based on the combined results of the first three regressions. Specifically, while both Perceived Fairness and Expected Disappointment are important factors in explaining participants' effort levels, the effect of Perceived Fairness is dominated by the more powerful opposing effect of Expected Disappointment.[8] We also note that a marginally significant effect of Contract Frame on Effort ($t = -1.95$, $p = .056$) remains after controlling for both Perceived Fairness and Expected Disappointment in the fourth regression, suggesting that, despite the high adjusted R^2 (.48), there is still some other force not fully captured in our Expected Disappointment measure that caused participants to choose more effort under the penalty contract than under the bonus contract.[9]

5. DISCUSSION

Our findings can be summarized as follows: Consistent with Luft (1994), we find that employees generally prefer bonus contracts to economically equivalent penalty contracts. We extend Luft's study by demonstrating that employee effort is higher under a penalty contract than an economically equivalent bonus contract. Our analysis demonstrates that this finding is the result of two effects that work in opposite directions, where the first effect dominates the second. The first effect is due to loss aversion, which makes employees more averse to having to pay a penalty than to not receiving a bonus, and causes them to choose more effort under the penalty contract. The second effect reflects reciprocity, which causes employees who consider their contracts to be fairer to choose more effort. Because employees generally perceived the bonus contract to be fairer than the penalty contract, reciprocity caused employees to choose more effort under the bonus contract. We find support for both of these opposing effects, with reciprocity dampening, but not completely offsetting, the dominant effect of loss aversion on employee effort.

Our results have important implications for understanding why in practice most actual contracts are framed as bonus contracts. Under conventional economic analysis, it is irrelevant whether a contract is framed in bonus terms or penalty terms. However, Luft argued that because employees prefer bonus contracts, they would demand higher payments from firms to accept penalty contracts. Thus, assuming no other differential cost between offering a bonus or penalty contract, firms maximize profits by offering bonus contracts. However, our results show that firms do bear a cost for offering bonus contracts rather than penalty contracts. That is, employees choose more effort under penalty contracts, so offering a bonus contract gives up the benefit of this increased effort. Consequently, it is no longer clear that offering a bonus contract maximizes firm profit.

Of course, this brings us back to the original question of why in practice most actual contracts are bonus contracts. Conventional economic theory offers no explanation because it predicts indifference between bonus contracts and economically equivalent penalty contracts. Given our results, Luft's (1994) explanation that we observe bonus contracts in the world because employees prefer them cannot be a complete explanation because the attendant logic fails to recognize the cost of forgone effort associated with bonus contracts. In our view, a better and more comprehensive explanation is likely to be found only through additional research designed to understand the full range of costs and benefits associated with each type of contract.

Along with Luft's (1994) study, the results of this study clearly show that conventional economic analysis fails to capture either employees' preferences for bonus contracts or the fact that penalty contracts motivate higher effort. Thus, we already know that there are important costs and benefits associated with incentive contracts beyond those typically assumed to exist in conventional economic analysis. We suspect there are still other costs associated with offering penalty contracts, or benefits associated with offering bonus contracts, that are not yet reflected in any of the currently available explanations for why penalty contracts are rarely observed in practice. For example, it may be that employees have ways to retaliate against a firm who offers them a penalty contract other than to withhold effort. In our experiment, withholding effort was the only possible means of retaliation against the firm. However, withholding effort also increased the chance of having to pay the penalty, and therefore employees who were motivated by loss aversion to avoid the penalty had no choice but to bear the higher cost of choosing higher effort.

In contrast, in some actual work settings, employees could choose to work hard to avoid the penalty just as they did in our experiment, but then quit their job or take other retaliatory actions against the firm that would benefit the employee at the expense of the firm. If firms anticipate that such employee retaliation will be a likely consequence of using penalty contracts, a cost/benefit analysis may lead them to conclude that using bonus contracts is likely to maximize profits in the long run.

Alternatively, in settings where employees have other options for extracting monetary benefits (for example employee theft), employees may in fact withhold effort (i.e., not exert more effort under a penalty contract than a bonus contract), but then extract monetary benefits through other means to make up for the forgone incentive pay (Greenberg 1990). In this case, there would be an extra cost, but no offsetting benefit, associated with using a penalty contract. Consequently, firms would likely maximize profits by offering bonus contracts. Future research could investigate whether such other potential costs of using penalty contracts can help explain why bonus contracts are typically observed in practice.

ACKNOWLEDGMENT

We appreciate helpful comments from an anonymous reviewer, Jake Birnberg, Jim Boatsman, Larry Brown, Bryan Church, Harry Evans, Yuhchang Hwang, Kathryn Kadous, Steve Kaplan, Ed O'Donnell, Casey Rowe, Kristy Towry, Bill Waller and workshop participants at the Second Asian Conference on Experimental Business Research, the American Accounting Association annual meeting, Arizona State University, Emory University, and the University of Pittsburgh. We also thank Fred Jacobs for providing experimental participants and Frank Luo for helping conduct the experiment. This research was supported by grants from the Robinson College of Business to Hannan and from the Katz Graduate School of Business to Hoffman and Moser.

NOTES

[1] A recent study by Frederickson and Waller (2004) supports Luft's (1994) findings in an experimental setting where firm participants and employee participants interact. Specifically, employees demanded higher expected pay to accept an offer framed as a penalty contract versus an offer framed as a bonus contract, and firms accommodated this employee loss aversion by making higher offers in the penalty contract frame. We note that Luft did not limit her explanation for why employees prefer bonus contracts to Kahneman and Tversky's (1979) pure prospect-theory definition of loss aversion. Nevertheless,

all the related effects she described are consistent with the more general notion of loss aversion which reflects the well-documented finding in psychology that the negative response to penalties, losses, punishment, etc. is greater than the positive response to equivalent bonuses, gains, rewards, etc. In addition, because all the loss-aversion related effects described by Luft (1994) and Frederickson and Waller (2004) cause employees to prefer bonus contacts to penalty contracts, they all lead to Luft's conclusion that employers maximize profit by offering bonus contracts.

[2] It was made clear to participants that they were required to pay back any penalty amounts before leaving the room. An additional experimenter was posted at the door to ensure that no participant exited the room before settling their accounts with the experiment's paymaster, and no participant attempted to do so.

[3] This operationalization of effort is consistent with the agency theory definition of effort (Baiman 1982) in that (1) participants had control over the effort level choice, (2) higher effort increased the probability of reaching the target (high) outcome, and (3) participants derived disutility from choosing a higher effort level.

[4] For example, choosing effort level 1 yielded an expected payoff of .30 ($30) + .70 ($20) − $0.50 = $22.50, while choosing effort level 10 yielded an expected payoff of .75 ($30) + .25 ($20) − $5 = $22.50.

[5] The total number of chips in each of the 13 bags varied as necessary to create the exact outcome probability distribution corresponding to each effort level.

[6] We also verified that the relation between Expected Disappointment and Effort was in the predicted direction, and that the path was not the opposite one. That is, it is not the case that greater effort led to greater disappointment. This additional test of the causal path is reported in footnote 8.

[7] Perceived Fairness and Expected Disappointment are not significantly correlated with each other [Pearson (Spearman) correlation = −.105 (−.087), $p > .40$ (.47)].

[8] We also used structural equation modeling to test the overall model depicted in Figure 1. Under this approach, a single covariance matrix is created and used to simultaneously estimate all links in the model (Kline 1998). Results from the overall model are almost identical to the corresponding individual tests reported for H2 – H5 and RQ1 and RQ2. Specifically, Expected Disappointment was higher under the Penalty contract and Perceived Fairness was higher under the Bonus contract. Both Expected Disappointment and Perceived Fairness had a positive effect on Effort. The path coefficients are equal to those reported in regression model 4 in Table 3 for Expected Disappointment to Effort (.62), Perceived Fairness to Effort (.48) and Contract Frame to Effort (−1.57). The primary measure of fit for the structural equation model is a Chi-square statistic, which tests the null hypothesis that the proposed model is a good fit for the data. For the model depicted in Figure 1, the Chi-square is not statistically significant ($\chi^2 = .000$, $p = .995$), indicating that the model is a good fit. To bolster confidence in our interpretation of the causal path, an alternative model was specified which reversed the direction between Expected Disappointment and Effort. That is, although we hypothesized that Expected Disappointment would result in higher Effort, a potential alternative explanation is that choosing higher effort caused participants to expect to be more disappointed if they had to pay the penalty or forgo the bonus. The null hypothesis that the model is a good fit for the data is rejected for the alternative model ($\chi^2 = 11.06$, $p < .005$), indicating that the alternative causal explanation is not tenable. A similar model which reversed the causal direction for Perceived Fairness also was not tenable ($\chi^2 = 23.14$, $p < .001$).

[9] We also ran a regression including all three independent variables included in the fourth regression and all related interactions ($R^2 = .60$). Results are consistent with those reported for the fourth regression in that both Expected Disappointment and Perceived Fairness remain statistically significant ($ps < .001$). In addition, neither the 3-way interaction, nor the 2-way interactions between Contract Frame and Perceived Fairness or Contract Frame and Expected Disappointment are statistically significant. The significant interaction between Perceived Fairness and Expected Disappointment ($t = -3.76$, $p < .01$) reflects a diminishing return to Effort when both Expected Disappointment and Perceived Fairness are high, and as such, has no implications for the interpretation of the results reported in the paper. Finally, the effect of Contract Frame drops to nonsignificance when all interactions are included in the model, suggesting that the remaining negative effect of Contract Frame on Effort noted in the fourth regression may reflect the fact that the model excluded the interaction terms.

REFERENCES

Baker, G. P., M. C. Jensen, and K. J. Murphy. (1988). "Compensation and incentives: Practice vs. theory." *The Journal of Finance*. 43(3): 593–616.

Baiman, S. (1982). "Agency research in managerial accounting: A survey." *Journal of Accounting Literature* 1 (Spring): 161–213.

Charness, C. and M. Rabin. (2002). "Understanding social preferences with simple tests." *The Quarterly Journal of Economics*, 117(3), 817–869.

Demski, J. S. and G. A. Feltham. (1978). "Economic incentives in budgetary control systems." *The Accounting Review*, 53(2), 336–359.

Fehr, E., S. Gächter, and G. Kirchsteiger. (1997). "Reciprocity as a contract enforcement device: Experimental evidence." *Econometrica*, 65(4), 833–860.

Fehr, E., E. Kirchler, A. Weichbold, and S. Gächter. (1998). "When social norms overpower competition: Gift exchange in experimental labor markets." *Journal of Labor Economics*, 16(2), 321–354.

Fehr, E., G. Kirchsteiger, and A. Riedl. (1993). "Does fairness prevent market clearing? An experimental investigation." *The Quarterly Journal of Economics*, 108(2), 438–459.

Feltham, G. and J. Xie. (1994). "Performance measure congruity and diversity in multi-task principal/agent relations." *The Accounting Review*, 69(3), 429–453.

Frederickson, J. R. (1992). "Relative performance information: The effects of common uncertainty and contract type on agent effort." *The Accounting Review*, 67(4), 647–669.

Frederickson, J. R. and W. Waller. (2004). Carrot or stick? Contract framing and the use of decision-influencing information in a principal–agent setting. Working paper, Hong Kong University of Science and Technology.

Goranson, R. E. and L. Berkowitz. (1966). "Reciprocity and responsibility reactions to prior help." *Journal of Personality and Social Psychology*, 3, 227–232.

Greenberg, J. (1978). "Effects of reward value and retaliative power on allocation decisions: Justice, generosity or greed?" *Journal of Personality and Social Psychology*, 36, 367–379.

Greenberg, J. (1990). "Employee theft as a reaction to underpayment inequity: The hidden cost of pay cuts." *Journal of Applied Psychology*, 75(5), 561–569.

Greenberg, M. S. and D. Frisch. (1972). "Effect of intentionality on willingness to reciprocate a favor." *Journal of Experimental Social Psychology*, 8, 99–111.

Hannan, R. L., J. H. Kagel, and D. V. Moser. (2002). "Partial gift exchange in experimental labor markets: Impact of subject population differences, productivity differences and effort requests on behavior." *Journal of Labor Economics,* 20(4), 923–951.

Holmstrom, B. (1979). "Moral hazard and observability." *Bell Journal of Economics*, 10(1), 74–91.

Holmstrom, B. (1982). "Moral hazard in teams." *Bell Journal of Economics*, 13(2), 324–340.

Holmstrom, B. and P. Milgrom. (1991). "Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design." *Journal of Law, Economics, & Organization*, 7, 24–52.

Kahneman, D., J. L. Knetsch, and R. H. Thaler. (1986). "Fairness as a constraint on profit seeking: entitlements in the market." *American Economic Review*, 76(4), 728–741.

Kahneman, D. and A. Tversky. (1979). "Prospect theory: An analysis of decision making under risk." *Econometrica*, 47, 263–291.

Kline, R. B. (1998). *Principles and Practice of Structural Equation Modeling*. New York, NY: The Guilford Press.

Luft, J. (1994). "Bonus and penalty incentives: Contract choice by employees," *Journal of Accounting and Economics*, 18(2), 181–206.

Milgrom, P. and J. Roberts. (1992). *Economics, Organizations and Management*. Englewood Cliffs: Prentice Hall.

Rabin, M. (1993). "Incorporating fairness into game theory and economics." *American Economic Review*, 83(5), 1281–1302.

Young, S. M. and B. Lewis. (1995). "Experimental incentive-contracting research in management accounting." In R. H. Ashton and A. H. Ashton, eds. *Judgment and Decision-Making Research in Accounting and Auditing*. U.K.: Cambridge University Press: 55–75.

Chapter 9

MANAGERIAL INCENTIVES AND COMPETITION

Rachel Croson
University of Pennsylvania

Arie Schinnar
University of Pennsylvania

Abstract

This paper experimentally tests the impact of managerial incentives on competitive (market) outcomes. We use a Cournot duopoly game to show that when managers' incentives are based on the firm's *absolute* performance (profits), collusion can be sustained. However, when managers' incentives are based on the firm's *relative* performance (their profits relative to the other firm's profits), this drives the market to the competitive and efficient outcome. These results suggest that regulators need to consider not only the number and concentration of firms in an industry, but also the managerial compensation schemes when deciding how much intervention is appropriate in a given industry.

1. INTRODUCTION

The question of how to motivate managers in order to maximize firm's profit is an important one, both economically and psychologically. A large literature examines the impact of managerial incentives on effort and performance, especially under conditions of moral hazard, where there exist principal-agent problems. A nice discussion and review of this issue can be found in Prendergast (1999).

In contrast, we are interested in the impact of managerial incentives on collusive outcomes of firms. Thus the experiment reported here builds on the work of Vickers (1985), Sklivas (1987) and Fershtman and Judd (1987). In these papers, the authors theoretically explore the question of which incentive schemes firms might choose for their managers, comparing those which reward managers only on profits with those that reward managers on some combination of profits and revenue (Sklivas) or profits and sales volume (Vickers, Fershtman and Judd).

In this paper, we will compare slightly different compensation schemes; one in which managers are paid as a function of the profits of the firm, and a second where

A. Rapoport and R. Zwick (eds.), Experimental Business Research, Vol. II, 171–184.

they are compensated based on their performance relative to the other firm in their industry. Unlike the previous papers, we will focus not on the incentives for the firms to choose one of these compensation schemes over another, but instead experimentally explore how managers act when compensated by each one.

We use a symmetric Cournot duopoly setting with perfect information and no uncertainty. When managers are compensated based on firm profits, the equilibrium of the game involves collusion. However, when managers are compensated based on relative profits, the equilibrium devolves to the perfectly competitive outcome. We test this simple theory in an experiment. Participants play a series of one-shot Cournot games in a strangers design. We find, consistent with the theory, that individuals produce significantly less quantity (are more collusive) when they are compensated based on their absolute performance than when they are compensated based on their relative performance.

These results are useful on a number of dimensions. First, they provide psychological support for the theory and its predictions. Second, they highlight the importance of firms' choice of managerial incentives to maximize own profit. Finally, they highlight an additional tool that regulators may have in preventing collusion – they can monitor executive compensation in addition to (or perhaps instead of) output in markets where collusion is suspected.

The remainder of this chapter is organized as follows. Section II introduces the Cournot setting and derives predictions using the parameters from the experiment. Section III presents the experimental design and implementation that we used. Section IV describes our results and section V concludes.

2. COURNOT COMPETITION

This model incorporates the basic Cournot intuition. Two managers work for symmetric firms and face a known (and here, linear) demand function. Each faces marginal costs (here, constant) and independently chooses the quantity their firm will produce. We assume that the manager chooses the quantity his firm will produce so as to maximize his own earnings, given his compensation package. We examine two cases, first, the case in which each manager is compensated with a fraction of his firm's absolute profits and second, the case in which each manager is compensated based on his profits relative to the other firm.

In the experiment, we use the following parameters:

$$\text{Demand function: Price} = \$100 - (q_1 + q_2)$$
$$\text{Marginal cost: } \$10 \text{ per unit}$$

Case A. The manager is compensated based on a fraction of his firm's profits; in the experiment he earns the firm's profits divided by \$1000. Thus for each \$1000 of firm profits, he earns \$1. Manager i thus maximizes his earnings by solving the following problem.

$$\max_{q_i} \{q_i * [(100 - q_i - q_j) - 10]\} / 1000.$$

Taking the other firm's production as given, this yields

$$q_i^* = q_j^* = 30,$$

the classic Cournot equilibrium outcome.[1] Note that this level of production yields $900 of firm profits, and thus 90 cents of profit for each manager in the experiment.[2]

Case B. Here, the manager is compensated based not on his absolute profit but based on his profit relative to the other manager. In our experiment, we implement this in the following way.

- If both firms earn the same amount of profit, each earns $1 in managerial compensation.
- If one firm earns more profit than another, the manager of the more profitable firm earns $2 in managerial compensation, and the manager of the less profitable firm earns $0 in managerial compensation.[3]

Note first that the Cournot outcome above is no longer an equilibrium. If each firm produces 30 units, then each firm earns $900 of profit and each manager earns $1 (since both firms have the same profit). However, a given manager can profitably deviate from this outcome, producing 31 units of output, decreasing his firm's profit to $899 but decreasing the profit of the other firm to $870. Since his firm is now more profitable than the other firm, he earns $2 of managerial compensation while the other manager earns $0.

This process continues until both firms are producing the perfectly competitive output of 45 units, charging a price equal to marginal cost of $10 and earning zero profits. Each manager thus earns $1. At this point, no manager wants to increase their production further. Increasing a firm's production to 46 units results in a profit of −46 for this firm and only −45 for the competitor, thus no manager wants to increase (or decrease) their production from this point.[4]

Two previous experiments have tested behavior in settings related to this one. In Potters, Rockenbach and Sadrieh (2004) the authors run an experiment in which managers are compensated based on relative productivity. However, their focus is on the extent to which the agents (managers) collude to take advantage of the principals (firms) when compensated in this way. In our relative compensation treatment, there is no benefit from such collusion; the total to be paid to managers is fixed. A second paper, Huck, Muller and Normann (forthcoming), directly tests the predictions of the Vickers/Fershtman and Judd models in which firms choose compensation schemes (either just profits or profits and sales) and managers choose quantities in response to those contracts. Our paper is a simpler version (examining just the managers' actions), and examines a different compensation scheme (relative profits rather than sales volume).

3. EXPERIMENTAL DESIGN AND IMPLEMENTATION

We test the predictions of this theory using an experiment involving Cournot competition. Many previous researchers have tested Cournot competition experimentally (see, for example, Holt 1995 for a review), starting with Fouraker and Siegel (1963). These experiments have found that participants often play Cournot equilibria when the number of competitors is sufficiently small (2 and sometimes 3) and when they have the opportunity to learn the game (repeated play).

We want an environment in which collusion occurs in the baseline case so as to show we can make it disappear when managers are compensated based on relative profits. Thus we will use the duopoly setting in our experiment with perfect and stationary information about the demand function, both competitor's (constant) marginal costs, and a repeated game setting.

Our experiment involved 91 participants: 43 in the baseline treatment and 48 in the relative payment treatment. Participants were recruited from the undergraduate population at the University of Pennsylvania, and were told that they could participate in an experiment in which they would keep their earnings. The experiment was run by hand, in a large classroom at the University. Participants were seated sufficiently far apart that they could not see the decisions of others.

Participants were brought into the lab and given a $5 show-up fee. Instructions were handed out and read aloud. A copy of the instructions can be found in Appendix A. The instructions included details of the participants' compensation schemes which differed between the treatments. Each participant represented a manager whose firm competes in a duopolistic market, producing goods that are perfect substitutes. Participants were then randomly matched, and were asked to make a production (quantity) decision in the Cournot setting. Each manager could choose the quantity they wished their firm to produce, and that choice in combination with their counterpart's choice determined their and their counterpart's profits. The parameters for the experiment were as described above.

Participants played the game eight or nine times, depending on the size of the group that had arrived in the lab. For each iteration, they were paired against a different counterpart (strangers design) with no overlap. No individual met the same counterpart more than once in the experiment. After each iteration, the participants were told their firm's profits, their counterpart's firm's profits, and their own earnings.[5] In the baseline treatment, participants were paid based on their firm's absolute profits earned over the entire experiment. In the relative payment treatment, participants' earnings were determined by comparing their firm's profits relative to their counterpart's profits in each round (as described above), and summed over the rounds that they played. Average earnings in the experiment were $12.63 in the baseline treatment and \$13.72 in the relative compensation treatment, including the show-up fee, and the experiment lasted about 45 minutes.

4. RESULTS

We compare behavior in this experiment to the two equilibrium predictions. In the baseline treatment, we predict that participants will choose the Cournot quantity (30). In the relative compensation treatment, we predict that participants will choose the competitive quantity (45). Thus we predict that quantities chosen will be higher when managers are compensated based on relative profits than when they are compensated based on absolute profits.

These predictions were supported by the data. The average quantity produced in the baseline treatment was 34.3, only slightly higher than the equilibrium prediction of 30. The average quantity produced in the relative compensation treatment as 43.6, only slightly lower than the equilibrium prediction of 45.

Since we have multiple observations for each participant, we first calculate each person's average production over the rounds he played. Figure 1 graphs the distribution of those productions in the two treatments, arranged in increasing order. A t-test comparing these distributions finds a significant difference between them ($p < .0001$).

We also report the results of an OLS regression of quantities chosen, using each period's observation. We control for individual effects, session effects, and the period number (repetition). Results are shown in Table 1, above.

There is a significant difference between the two treatments; in particular, the quantity produced is significantly lower in the baseline treatment than in the relative compensation treatment. In addition, we also see a small but significant coefficient on the period number. Quantities tend to increase over time, suggesting that participants may be trying to collude in early rounds, but converging toward equilibrium once they experience the game.[6]

While comparative-statics of this experiment support the theory's predictions, the levels are not quite as predicted. That is, while managers produce more under relative compensation than under absolute compensation (the baseline), they don't produce *enough* more. The coefficient on treatment in the above regression is significantly greater than the predicted coefficient of -15 ($p < .01$). Thus we interpret our results as qualified support for the theory.

A further analysis investigates the time-trend of production decisions. Although each period represents a one-shot game, participants in the experiment may be learning how to play, and behavior may converge toward or away from the equilibrium. Figure 2 graphs average quantities produced in each round of the game in the two treatments. Note the axes in this figure; the minimum is 30, which is the predicted Cournot production. The maximum is 45, which is the predicted relative production.

As can be seen, there is a definite trend of increasing quantities in the relative payment treatment. The trend in the baseline (Cournot) treatment is less clear.

Table 1. OLS Regression of Quantity Chosen

	Estimate	*t-statistic*	*p-value*
Intercept	36.78	66.11	0.0000
Baseline Treatment	−4.65	18.20	0.0001
Period	0.444	4.46	0.0001
Session dummies		yes	
Individual dummies		yes	
	N R^2 (adj.)		758 .4765

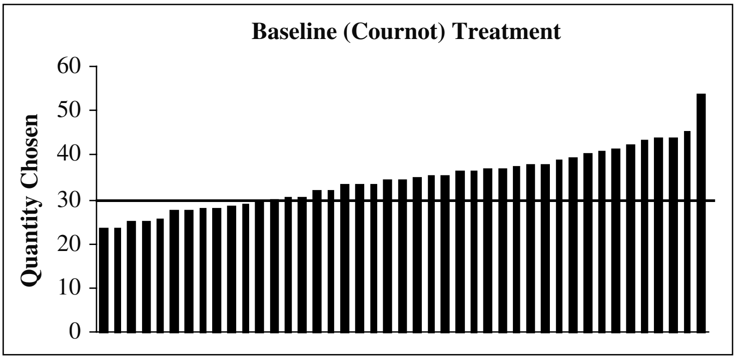

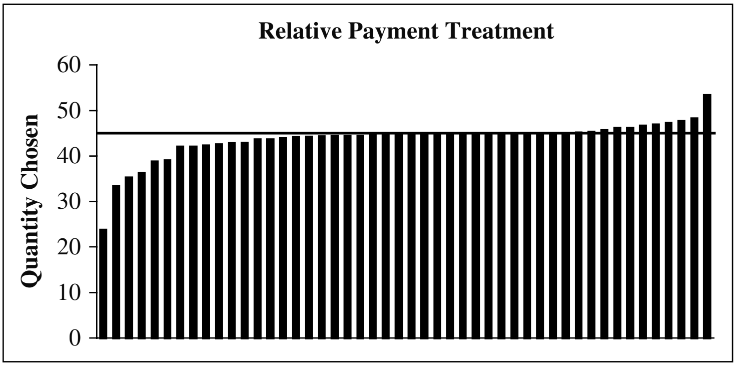

Figure 1. Each Participant's Average Quantities Chosen.

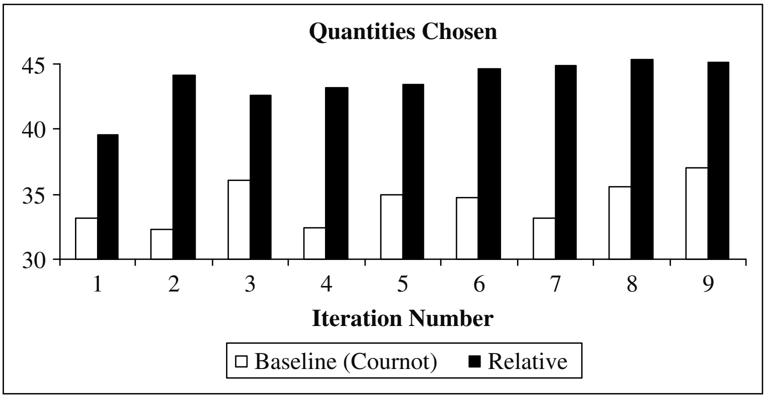

Figure 2. Average Quantities Chosen Over Time.

Table 2. OLS Regression of Quantity Chosen

	Baseline (Cournot)			*Relative Payment*		
	Estimate	*t-statistic*	*p-value*	*Estimate*	*t-statistic*	*p-value*
Intercept	32.44	35.96	0.0000	40.98	62.20	0.0000
Period	0.180	1.13	0.2596	0.628	4.51	0.0000
Session dummies		yes			yes	
Individual dummies		yes			yes	
	N 344 R^2 (adj.) .3884			N 414 R^2 (adj.) .2938		

Indeed, separate regressions for each treatment show a significant time-trend in the former, but not in the latter.

5. CONCLUSIONS AND IMPLICATIONS

Our experimental results support the model's comparative-static predictions: how managers are compensated (based on absolute or relative profits) has important implications for collusive behavior. Further research might investigate individual behavior in this setting by looking at each participant's history and seeing how they react to particular outcomes. Future research might also expand the scope of the inquiry to include the decision(s) of the firms in choosing these contracts.

In any experiment proporting to say something about the world, external validity questions are of extreme importance. One limitation of this analysis and experiment is that it applies only to firms whose products are perfect substitutes (are not horizontally or vertically differentiated). That said, a number of other asymmetries can be added to the model without changing its results, including asymmetric (but constant) marginal costs, asymmetric proportions of absolute profits that managers are compensated, and asymmetric measures of relative profitability.

In practice, many compensation contracts involve fixed salaries and are not based on beating the competition. However, there may be other reasons why managers would care about relative profits. For example, if part of managerial compensation is through stock and options, and the market incorporates relative performance into its valuations, then managers' compensation will include relative performance. Furthermore, relative performance may influence whether a manager keeps his job or gets promoted.

In addition to validating the theory, these results have important lessons for antitrust regulators. To determine whether an industry is collusive it is not sufficient (and may not even be necessary) to look at the industry's output, one should also look at managerial incentives of the individual firms. Similarly, regulating managerial incentives may have a bigger impact than simply denying specific mergers. Even in very concentrated (two-party) industries in our experiment, when incentives were relative rather than absolute, outcomes were competitive. Thus even in industries where concentration and other usual measures of collusive potential are the same, the amount of inefficiency that is observed is likely to depend on the incentives of the managers.

This research also raises an additional cost to firms compensating their managers based on relative performance. On the one hand, these compensation schemes can overcome principal/agent problems when there exist informational asymmetries. On the other hand, they may lead to incentives which reduce firms' profits. Firms need to balance the costs and benefits when considering varying compensation schemes.

NOTES

[1] We have calculated this solution for the case of symmetric firms. But it is robust to asymmetries in constant marginal cost, in the exact proportion of profits which each manager is compensation (whether they are the same or different), and size of the firm (capacity) so long as that constraint does not bind. One assumption about Cournot competition on which we rely quite heavily is the assumption of homogeneous products. If the firms are producing products which are vertically or horizontally differentiated, then Cournot is no longer the appropriate model.

[2] One concern with Cournot experiments like this one is the *flat maximum critique* (Harrison 1989), which argues that payoffs are not sensitive to participants' actions around equilibrium. That is, if my partner is producing the Cournot equilibrium quantity, the costs to me from deviating from the equilibrium are quite low. This can add noise to the outcomes. As will be seen, the second treatment (case B, relative compensation), does not suffer from this critique; there a deviation is quite costly. This critique predicts that the variance of quantities chosen in the first treatment will be higher than those chosen in the second treatment. This prediction is in fact true; the standard deviation of quantities chosen in the first treatment is 9.76 and in the second is 7.31. An F-test suggests that this difference is significant, $p < .01$.

[3] The reasonableness of this relative payment scheme relies on the symmetry of the two firms. If instead the firms vary, for example, on marginal cost, then comparing absolute profits is clearly not a reasonable benchmark. However, if managers are compensated relatively but based on standardized benchmarks (e.g., return on capital, profits relative to size, . . .), the same competitive results hold.

[4] Others have used experimental designs in which participants are compensated based on their relative payoffs, but in different contexts and for different purposes. Andreoni (1995) and Kurzban and Hauser (2002) use relative payoffs in public goods games to differentiate kindness and confusion. Croson and Donohue (2003), Croson and Donohue (forthcoming) use relative payoffs to capture real-world benchmarking incentives in a supply chain management game.

[5] In addition, we ran a third treatment in which 31 participants played a Cournot game similar to our baseline treatment. However, after each round of the game they were not told the profits of their counterpart. Theoretically this should not make a difference, and indeed empirically there were no differences between this treatment of incomplete information and the Cournot (baseline) treatment we report here. Instructions and data from this treatment are available from the authors.

[6] A similar regression using a discrete measure of period (dummies for periods 2 through 9) yields identical results. We also look for and fail to find an interaction between the treatment and the period number; thus any observed learning is occurring at the same speed in the two treatments.

REFERENCES

Andreoni, James (1995). "Cooperation in Public-Goods Experiments: Kindness or Confusion?" *American Economic Review*, 85(4), 891–904.

Croson, Rachel and Karen Donohue (2003). "The Impact of POS Data Sharing on Supply Chain Management: An Experimental Study." *Production and Operations Management*, 12, 1–11.

Croson, Rachel and Karen Donohue (forthcoming). Behavioral Causes of the Bullwhip Effect and the Observed Value of Inventory Information. Management Science.

Fershtman, Chaim and Kenneth Judd (1987). "Equilibrium Incentives in Oligopoly." *American Economic Review*, 77, 927–940.

Fouraker, Lawrence and Sidney Siegel (1963). Bargaining Behavior. McGraw-Hill: New York.

Harrison, Glenn (1989). "Theory and Misbehavior of First Price Auctions." *American Economic Review*, 79, 749–62.

Houser, Daniel and Robert Kurzban (2002). "Revisiting kindness and confusion in public goods games." *The American Economic Review*, 92(4), 1062–1069.

Holt, Charles (1995). Industrial Organization. In Handbook of Experimental Economics (Kagel and Roth, eds.), Princeton University Press: Princeton, NJ. 349–444.

Huck, Steffen, Wieland Muller and Hans-Theo Norman (forthcoming). Strategic Delegation in Experimental Markets. International Journal of Industrial Organization.

Potters, Jan, Bettina Rockenbach and Abdolkarim Sadrieh (2004). Collusion Under Yardstick Competition: An Experimental Study. Working Paper, Tilburg University.

Prendergast, Candice (1999). "The Provision of Incentives in Firms." *Journal of Economic Literature*, XXXVII, 7–63.

Sklivas, Steven (1987). "The Strategic Choice of Managerial Incentives." *RAND Journal of Economics*, 18(3), 452–458.

Vickers, John (1985). "Delegation and the Theory of the Firm." *The Economic Journal*, 95, 138–147.

APPENDIX A: INSTRUCTIONS

CC

Instructions

Welcome!

In this session you will be playing the role of a firm competing in a marketplace. You and another student representing a competing firm will simultaneously decide how much of a commodity product to produce. The quantities you choose will be combined to determine the price at which you can sell your product, and your corresponding profit. At the end of the session, you will receive cash earnings corresponding to your firm's profitability. The more you earn as a firm, the more money you as an individual will earn.

Your Firm

Imagine that you and another competing firm both manufacture an identical product called a widget. There are no fixed costs of production, but each widget you manufacture costs you $10. Your competitor has the same costs as you do.

The Marketplace

In this market, you and your competitor simultaneously choose how many widgets to manufacture, incurring the cost of $10 for each widget produced. Given the total quantity produced by you and your competitor, the market determines a price which it will pay for your widgets according to the following formula:

$$\text{Price} = 100 - [\text{your quantity} + \text{competitor's quantity}]$$

Your revenue from this market would thus be the price as above, times the number of widgets you produced. Your costs would be $10 times the number of widgets you produced. Your profit would be your revenue minus your costs. Notice that it is possible to earn negative profit in this market (if you and your competitor together produce more than 100 widgets).

The Interaction

We want to let you have some practice in making this decision, thus in this session you will be playing this role multiple times. However, each time you make a decision of how much to produce, you will be facing a **different** competitor. You will never compete against the same rival twice.

Everyone in the room has been assigned an ID number. At the beginning of each round you will complete a **Decision Form** which will tell you the ID number (but not the name) of your competitor. You will never face the same competitor more than once. After each round we will tell you your profit and the profit of your competitor for the round.

Earnings

At the end of the session we will add together your firm's profit in each round. You will earn money at the rate of one dollar for each $1000 of profit your firm has earned The more money your firm takes in profit, the more money you will earn.

Any questions?

Before we begin, let's look together at the forms you will be using for the experiment. At the beginning of each round you will see a form like the sample below.

Sample Decision Form

Round Number ___1___

Your ID Number ___1___ Your Name: ____________________

Your Competitor's ID Number ___2___

Your Quantity Produced ______

do not write below this line

Your Profit ______

Your Competitor's Profit ______

Earnings ______

This form will be used to keep track of your decisions and profit. The top of the form tells you the round number, your ID number and your competitor's ID number. This will be already filled in for you on each form, but you will have to fill in your name.

At the beginning of the round, you decide how any widgets to produce. Fill in that amount on the first line (Your Quantity Produced). Then hold the form above your head. The experimenter will come around and pick it up from you.

The experimenter will record your profit, the profit of your competitor and your earnings from this round (your profit ÷ 1000), and then hand the form back to you. Keep it to refer to at the end of the session. You may then proceed to the next round, decide how many widgets to produce and hold up the form for the experimenter to pick up. Remember, you will be facing different competitors in different rounds. You will not be matched with the same competitor twice.

At the end of the experiment, calculate your total cash earnings on the final form. We will call your name, bring all your forms with you to receive your cash earnings.

Any questions before we begin?

E

Instructions

Welcome!

In this session you will be playing the role of a firm competing in a marketplace. You and another student representing a competing firm will simultaneously decide how much of a commodity product to produce. The quantities you choose will be combined to determine the price at which you can sell your product, and your corresponding profit. At the end of the session, you will receive cash earnings corresponding to your firm's profitability relative to your competitor. The more you earn as a firm, relative to the competitor in your industry, the more money you as an individual will earn.

Your Firm

Imagine that you and another competing firm both manufacture an identical product called a widget. There are no fixed costs of production, but each widget you manufacture costs you \$10. Your competitor has the same costs as you do.

The Marketplace

In this market, you and your competitor simultaneously choose how many widgets to manufacture, incurring the cost of \$10 for each widget produced. Given the total quantity produced by you and your competitor, the market determines a price which it will pay for your widgets according to the following formula:

$$\text{Price} = 100 - [\text{your quantity} + \text{competitor's quantity}]$$

Your revenue from this market would thus be the price as above, times the number of widgets you produced. Your costs would be \$10 times the number of widgets you produced. Your profit would be your revenue minus your costs. Notice that it is possible to earn negative profit in this market (if you and your competitor together produce more than 100 widgets).

Next, we will compare the profit you earn with the profit earned by your competitor. If you earned more profit than your competitor, you will earn two dollars, if you earned less profit, you will earn zero dollars. If you and your competitor earned identical profit, you will each earn one dollar.

The Interaction

We want to let you have some practice in making this decision, thus in this session you will be playing this role multiple times. However, each time you make a decision of how much to produce, you will be facing a **different** competitor. You will never compete against the same rival twice.

Everyone in the room has been assigned an ID number. At the beginning of each round you will complete a **Decision Form** which will tell you the ID number (but not the name) of your competitor. You will never face the same competitor more than once. After each round we will tell you your profit and the profit of your competitor.

Earnings

Each round in which you earn strictly more profit than your competitor, you earn two dollars. Each round in which you and your competitor earn equal profit, you earn one dollar. At the end of the session we will add together the money you earned in each round. The more profit you make relative to your competitor, the more money you will earn.

Any questions?

Before we begin, let's look together at the forms you will be using for the experiment. At the beginning of each round you will see a form like the sample below.

Sample Decision Form

Round Number ___1___

Your ID Number ___1___ Your Name: ____________________

Your Competitor's ID Number ___2___

Your Quantity Produced ______

do not write below this line

Your Profit ______

Your Competitor's Profit ______

Earnings ______

This form will be used to keep track of your decisions and profit. The top of the form tells you the round number, your ID number and your competitor's ID number. This will be already filled in for you on each form, but you will have to fill in your name.

At the beginning of the round, you decide how any widgets to produce. Fill in that amount on the first line (Your Quantity Produced). Then hold the form above your head. The experimenter will come around and pick it up from you.

The experimenter will record your profit, the profit of your competitor and your earnings from the round, and then hand the form back to you. Keep it to refer to at the end of the session. You may then proceed to the next round, decide how many widgets to produce and hold up the form for the experimenter to pick up. Remember, you will be facing different competitors in different rounds. You will not be matched with the same competitor twice.

At the end of the experiment calculate your total cash earnings on the final form. We will call your name, bring all your forms with you to receive your cash earnings.

Any questions before we begin?

Chapter 10

DYNAMIC STABILITY OF NASH-EFFICIENT PUBLIC GOODS MECHANISMS: RECONCILING THEORY AND EXPERIMENTS

Yan Chen
University of Michigan

Abstract

We propose to use supermodularity as a robust dynamic stability criterion for public goods mechanisms with a unique Nash equilibrium. Among existing public goods mechanisms whose Nash equilibria are Pareto efficient, the Groves-Ledyard mechanism is a supermodular game if and only if the punishment parameter is sufficiently high, while none of the Hurwicz, Walker and Kim mechanisms is supermodular in a quasilinear environment. The Falkinger mechanism is a supermodular game in a quadratic environment if and only if the subsidy coefficient is greater than or equal to one. These results are consistent with the findings in seven experimental studies.

Keywords: public goods mechanisms, supermodular games, experiments
JEL Classification: H41, C62, D83

1. INTRODUCTION

The design of decentralized institutions to provide public goods has been a challenging problem for economists for a long time. Since the 1970s, economists have been seeking informationally decentralized mechanisms (i.e., mechanisms which only use the agents' own messages) that are non-manipulable (or dominant strategy incentive-compatible) and achieve a Pareto optimal allocation of resources with public goods. Some mechanisms have been discovered which have the property that Nash equilibria are Pareto optimal. These can be found in the work of Groves and Ledyard (1977), Hurwicz (1979), Walker (1981), Tian (1989), Kim (1993), Peleg (1996) and Falkinger (1996).

So far Nash implementation theory has mainly focused on establishing static properties of the equilibria. When a mechanism is implemented among real people, i.e., boundedly rational agents, however, we expect the actual implementation to be a dynamic process, starting somewhere off the equilibrium path. Following Hurwicz

A. Rapoport and R. Zwick (eds.), Experimental Business Research, Vol. II, 185–200.

(1972), one could interpret the Nash equilibrium strategies of a game form as the stationary messages of some decentralized learning process. The fundamental question concerning implementation of a specific mechanism is whether the dynamic processes will actually converge to one of the equilibria promised by theory. This paper addresses this question by proposing supermodularity as a robust stability criterion for public goods mechanisms when there is a unique Nash equilibrium.

The few theoretical papers on the dynamic properties of public goods mechanisms have been using very specific learning dynamics to investigate the stability of mechanisms. Muench and Walker (1983) and de Trenqualye (1988) study the convergence of the Groves-Ledyard mechanism under Cournot best-reply dynamics. De Trenqualye (1989) and Vega-Redondo (1989) propose mechanisms for which the Cournot best-reply dynamics is globally convergent to the Lindahl equilibrium[1] outcome. Kim (1993) proposed a mechanism which implements Lindahl allocations and remains stable under the gradient adjustment process given quasilinear utility functions. One exception is Cabrales (1999) who studies dynamic convergence and stability of the canonical mechanism in Nash implementation and the Abreu-Matsushima mechanism under "naive adaptive dynamics," which is different from the adaptive learning in Milgrom and Roberts (1990).

Recent experimental studies on learning strongly reject the Cournot best-reply learning model in favor of other models (e.g., Boylan and El-Gamal, 1993). So far there has been no experimental investigation of the gradient adjustment process, even though it has been used fairly extensively in the theoretical research on stability of games. Experimental research on learning is still far from reaching a conclusion with regard to a single "true" learning model that describes all adaptive behaviors. Furthermore, there is strong evidence that individual players adopt different learning rules under different circumstances (El-Gamal and Grether, 1995). It is therefore desirable to identify mechanisms which converge under a wide class of learning dynamics. This paper does so by focusing on mechanisms which are supermodular games.

Supermodular games (Milgrom and Roberts, 1990) are games in which the incremental return to any player from increasing her strategy is a nondecreasing function of the strategy choices of other players. Furthermore, if a player's strategy space has more than one dimension, components of a player's strategy are complements. Supermodular games encompass important economic applications of noncooperative game theory. For example, in games of new technology adoption, such as those in Dybvig and Spatt (1983), when more users hook into a communication system, the marginal return to others of doing the same often increases.

The class of supermodular games has been identified as having very robust dynamic stability properties (Milgrom and Roberts, 1990): it converges to the set of Nash equilibria that bound the serially undominated set under a wide class of interesting learning dynamics, including Bayesian learning, fictitious play, adaptive learning, Cournot best-reply and many others.[2] Therefore, instead of using a specific learning dynamic, we investigate whether we can find Nash-efficient public goods mechanisms which are supermodular games.

The idea of using supermodularity as a robust stability criterion for Nash-efficient mechanisms is not only based on its good theoretical properties, but also on strong experimental evidence. In fact it is inspired by the experimental results of Chen and Plott (1996) and Chen and Tang (1998), where they varied a punishment parameter in the Groves-Ledyard mechanism in a set of experiments and obtained totally different dynamic stability results.

In this paper, we review the main experimental findings on the dynamic stability of Nash-efficient public goods mechanisms, examine the supermodularity of existing Nash-efficient public goods mechanisms, and use the results to sort a class of experimental findings.

Section 2 introduces the environment. Section 3 reviews the experimental results. Section 4 discusses supermodular games. Section 5 investigates whether the existing mechanisms are supermodular games. Section 6 concludes the paper.

2. A PUBLIC GOODS ENVIRONMENT

We first introduce notation and the economic environment. Most of the experimental implementations of incentive-compatible mechanisms use a simple environment. Usually there is one private good x, one public good y, and $n \geq 3$ players, indexed by subscript i. Production technology for the public good exhibits constant returns to scale, i.e., the production function $f(\cdot)$ is given by $y = f(x) = x/b$ for some $b > 0$. Preferences are largely restricted to the class of quasilinear preferences, except Harstad and Marrese (1982) and Falkinger et al. (2000). Let E represent the set of transitive, complete and convex individual preference orderings, $\succcurlyeq_i$, and initial endowments, ω_i^x. We formally define E^Q as follows.

DEFINITION 1. $E^Q = \{(\succcurlyeq_i, \omega_i^x) \in E\colon \succcurlyeq_i$ is representable by a C^2 utility function of the form $v_i(y) + x_i$ such that $Dv_i(y) > 0$ and $D^2 v_i(y) < 0$ for all $y > 0$, and $\omega_i^x > 0\}$, where D^k is the k^{th} order derivative.

Falkinger et al. (2000) use a quadratic environment in their experimental study of the Falkinger mechanism. We define this environment as E^{QD}.

DEFINITION 2. $E^{QD} = \{(\succcurlyeq_i, \omega_i^x) \in E\colon \succcurlyeq_i$ is representable by a C^2 utility function of the form $A_i x_i - \frac{1}{2} B_i x_i^2 + y$ where $A_i, B_i > 0$ and $\omega_i^x > 0\}$.

An *economic mechanism* is defined as a non-cooperative game form played by the agents. The game is described in its normal form. In all mechanisms considered in this paper, the implementation concept used is Nash equilibrium. In the Nash implementation framework the agents are assumed to have complete information about the environment while the designer does not know anything about the environment.

3. EXPERIMENTAL RESULTS

Seven experiments have been conducted with mechanisms having Pareto-optimal Nash equilibria in public goods environments (see Chen (forthcoming) for a survey).

Sometimes the data converged quickly to the Nash equilibria; other times it did not. Smith (1979) studies a simplified version of the Groves-Ledyard mechanism which balanced the budget only in equilibrium. In the five-subject treatment (R1) one out of three sessions converged to the stage game Nash equilibrium. In the eight-subject treatment (R2) neither session converged to the Nash equilibrium prediction. Harstad and Marrese (1981) found that only three out of twelve sessions attained approximately Nash equilibrium outcomes under the simplified version of the Groves-Ledyard mechanism. Harstad and Marrese (1982) studied the complete version of the Groves-Ledyard mechanism in Cobb-Douglas economies. In the three-subject treatment one out of five sessions converged to the Nash equilibrium. In the four-subject treatment one out of four sessions converged to one of the Nash equilibria. Mori (1989) compares the performance of a Lindahl process with the Groves-Ledyard mechanism. He ran five sessions for each mechanism, with five subjects in each session. The aggregate levels of public goods provided in each of the Groves-Ledyard sessions were much closer to the Pareto optimal level than those provided using a Lindahl process. At the individual level, each of the five sessions stopped within ten rounds when every subject repeated the same messages. However, since individual messages must be in multiples of .25 while the equilibrium messages were not on the grid, convergence to Nash equilibrium messages was approximate. None of the above experiments studied the effects of the punishment parameter, which determines the magnitude of punishment if a player's contribution deviates from the mean of other players' contributions, on the performance of the mechanism.

Chen and Plott (1996) first assessed the performance of the Groves-Ledyard mechanism under different punishment parameters. Each group consisted of five players with different preferences. They found that by varying the punishment parameter the dynamics and stability changed dramatically. This finding was replicated by Chen and Tang (1998) with twenty-one independent sessions and a longer time series (100 rounds) in an experiment designed to study the learning dynamics. Chen and Tang (1998) also studied the Walker mechanism (Walker, 1981) in the same economic environment.

Figure 1 presents the time series data from Chen and Tang (1998) for two out of five types of players. The data for the remaining three types of players display very similar patterns. Each type differ in their marginal utility for the public good. Each graph presents the mean (the black dots), standard deviation (the error bars) and stage game equilibria (the dashed lines) for each of the two different types averaged over seven independent sessions for each mechanism. The two graphs in the first column display the mean contribution (and standard deviation) for types 1 and 2 players under the Walker mechanism (hereafter Walker). The second column displays the average contributions for types 1 and 2 for the Groves-Ledyard mechanism under a low punishment parameter (hereafter GL1). The third column displays the same information for the Groves-Ledyard mechanism under a high punishment parameter (hereafter GL100). From these graphs, it is apparent that all seven sessions of the Groves-Ledyard mechanism under a high punishment parameter converged[3] very quickly to its stage game Nash equilibrium and remained stable,

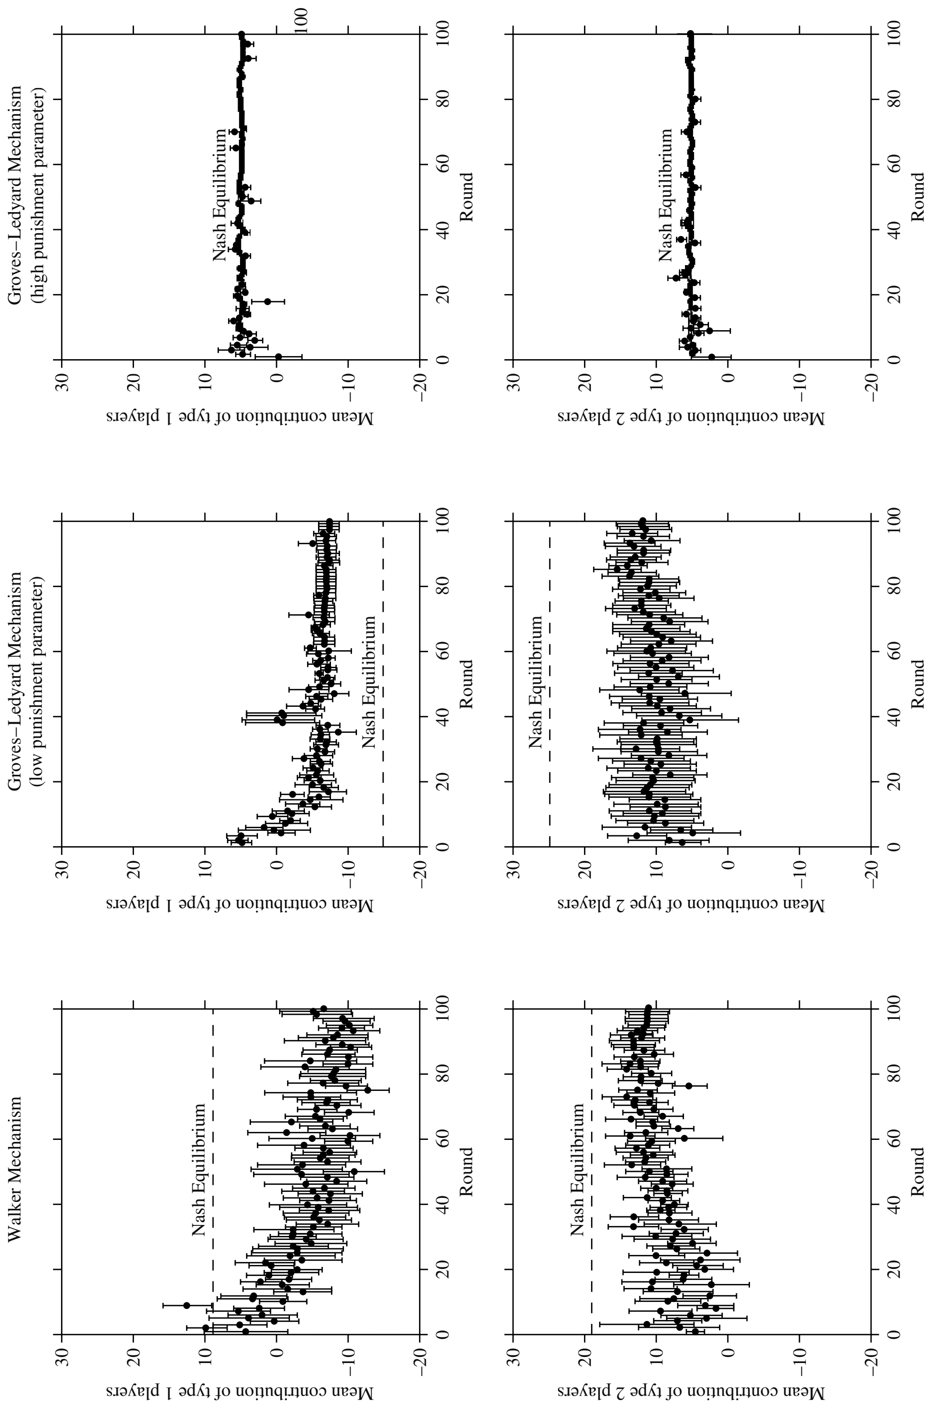

Figure 1. Mean Contribution and Standard Deviation in Chen and Tang (1998).

while the same mechanism did not converge under a low punishment parameter; the Walker mechanism did not converge to its stage game Nash equilibrium either.

Because of its good dynamic properties, GL100 had far better performance than GL1 and Walker, evaluated in terms of system efficiency, close to Pareto optimal level of public goods provision, less violations of individual rationality constraints and convergence to its stage game equilibrium. All these results are statistically highly significant (Chen and Tang, 1998).

These results illustrate the importance to design mechanisms which not only have good static properties, but also good dynamic stability properties like GL100. Only when the dynamics lead to the convergence to the static equilibrium, can all the nice static properties be realized.

Falkinger et al. (2000) study the Falkinger mechanism in a quasilinear as well as a quadratic environment. In the quasilinear environment, the mean contributions moved towards the Nash equilibrium level but did not quite reach the equilibrium. In the quadratic environment the mean contribution level hovered around the Nash equilibrium, even though none of the 23 sessions had a mean contribution level exactly equal to the Nash equilibrium level in the last five rounds. Therefore, Nash equilibrium was a good description of the average contribution pattern, although individual players did not necessarily play the equilibrium.

In Section 5 we will provide a theoretical explanation for the above experimental results in light of supermodular games.

4. SUPERMODULARITY AND STABILITY

We first define supermodular games and review their stability properties. Then we discuss alternative stability criteria and their relationship with supermodularity.

Supermodular games are games in which each player's marginal utility of increasing her strategy rises with increases in her rival's strategies, so that (roughly) the player's strategies are "strategic complements." Supermodular games need an order structure on strategy spaces, a weak continuity requirement on payoffs, and complementarity between components of a player's own strategies, in addition to the above-mentioned strategic complementarity between players' strategies. Suppose each player i's strategy set S_i is a subset of a finite-dimensional Euclidean space R^{k_i}. Then $S \equiv \times_{i=1}^{n} S_i$ is a subset of R^k, where $k = \sum_{i=1}^{n} k_i$.

DEFINITION 3. A *supermodular game* is such that, for each player i, S_i is a nonempty sublattice of R^{k_i}, u_i is upper semi-continuous in s_i for fixed s_{-i} and continuous in s_{-i} for fixed s_i, u_i has increasing differences in (s_i, s_{-i}), and u_i is supermodular in s_i.

Increasing differences says that an increase in the strategy of player i's rivals raises her marginal utility of playing a high strategy. The supermodularity assumption ensures complementarity among components of a player's own strategies. Note that it is automatically satisfied when S_i is one-dimensional. As the following theorem indicates supermodularity and increasing differences are easily characterized for smooth functions in R^n.

THEOREM 1. (*Topkis, 1978*) Let u_i be twice continuously differentiable on S_i. Then u_i has increasing differences in (s_i, s_j) if and only if $\partial^2 u_i / \partial s_{ih} \partial s_{jl} \geq 0$ for all $i \neq j$ and all $1 \leq h \leq k_i$ and all $1 \leq l \leq k_j$; and u_i is supermodular in s_i if and only if $\partial^2 u_i / \partial s_{ih} \partial s_{il} \geq 0$ for all i and all $1 \leq h < l \leq k_i$.

Supermodular games are of interest particularly because of their very robust stability properties. Milgrom and Roberts (1990) proved that in these games the set of learning algorithms consistent with adaptive learning converge to the set bounded by the largest and the smallest Nash equilibrium strategy profiles. Intuitively, a sequence is consistent with adaptive learning if players "eventually abandon strategies that perform consistently badly in the sense that there exists some other strategy that performs strictly and uniformly better against every combination of what the competitors have played in the not too distant past." (Milgrom and Roberts, 1990) This includes a wide class of interesting learning dynamics, such as Bayesian learning, fictitious play, adaptive learning, Cournot best-reply and many others.

Since experimental evidence suggests that individual players tend to adopt different learning rules (El-Gamal and Grether, 1995), instead of using a specific learning algorithm to study stability, one can use supermodularity as a robust stability criterion for games with a unique Nash equilibrium. For supermodular games with a unique Nash equilibrium, we expect any adaptive learning algorithm to converge to the unique Nash equilibrium, in particular, Cournot best-reply, fictitious play and adaptive learning. Compared with stability analysis using Cournot best-reply dynamics, supermodularity is much more robust and inclusive in the sense that it implies stability under Cournot best-reply and many other learning dynamics mentioned above.

5. SUPERMODULARITY OF EXISTING NASH-EFFICIENT PUBLIC GOODS MECHANISMS

In this section we investigate the supermodularity of five well-known Nash-efficient public goods mechanisms. We use supermodularity to analyze the experimental results on Nash-efficient public goods mechanisms.

The Groves-Ledyard mechanism (1977) is the first mechanism in a general equilibrium setting whose Nash equilibrium is Pareto optimal. The mechanism allocates private goods through the competitive markets and public goods through a government allocation-taxation scheme that depends on information communicated to the government by consumers regarding their preferences. Given the government scheme, consumers find it in their best interest to reveal their true preferences for public goods. The mechanism balances the budget both on and off the equilibrium path, but it does not implement Lindahl allocations. Later on, more game forms have been discovered which implement Lindahl allocations in Nash equilibrium. These include Hurwicz (1979), Walker (1981), Tian (1989), Kim (1993) and Peleg (1996).

DEFINITION 4. For the Groves-Ledyard mechanism, the strategy space of player i is $S_i \subset R^1$ with generic element $m_i \in S_i$. The outcome function of the public good and the net cost share of the private good for player i are

$$Y(m) = \sum_k m_k$$

$$T_i^{GL}(m) = \frac{Y(m)}{n} \cdot b + \frac{\gamma}{2}\left[\frac{n-1}{n}(m_i - \mu_{-i})^2 - \sigma_{-i}^2\right].$$

where $\gamma > 0$, $n \geq 3$, $\mu_{-i} = \sum_{j \neq i} m_j/(n-1)$ is the mean of others' messages, and $\sigma_{-i}^2 = \sum_{h \neq i}(m_h - \mu_{-i})^2/(n-2)$ is the squared standard error of the mean of others' messages.

In the Groves-Ledyard mechanism each agent reports m_i, the increment (or decrement) of the public good player i would like to add to (or subtract from) the amounts proposed by others. The planner sums up the individual contributions to get the total amount of public good, Y, and taxes each individual based on her own message, and the mean and sample variance of everyone else's messages. Thus each individual's tax share is composed of three parts: the per capita cost of production, $Y \cdot b/n$, plus a positive multiple, $\gamma/2$, of the difference between her own message and the mean of others' messages, $(n-1)/n \times (m_i - \mu_{-i})^2$, and the sample variance of others' messages, σ_{-i}^2. While the first two parts guarantee that Nash equilibria of the mechanism are Pareto optimal, the last part insures that budget is balanced both on and off the equilibrium path. Note that the free parameter, γ, determines the magnitude of punishment when an individual deviates from the mean of others' messages. It does not affect any of the static theoretical properties of the mechanism.

Chen and Plott (1996) and Chen and Tang (1998) found that the punishment parameter, γ, had a significant effect in inducing convergence and dynamic stability. For a large enough γ, the system converged to its stage game Nash equilibrium very quickly and remained stable; while under a small γ, the system did not converge to its stage game Nash equilibrium. In the following proposition, we provide a necessary and sufficient condition for the mechanism to be a supermodular game given quasilinear preferences, and thus to converge to its Nash equilibrium under a wide class of learning dynamics.

PROPOSITION 1. The Groves-Ledyard mechanism is a supermodular game for any $e \in E^Q$ if and only if $\gamma \in [-\min_{i \in N}\left\{\frac{\partial^2 v_i}{\partial y^2}\right\}n, +\infty]$.

Proof: Since u_i is C^2 on S_i, by Theorem 1, u_i has increasing differences in (m_i, m_{-i}) if and only if

$$\frac{\partial^2 u_i}{\partial m_i \partial m_j} = \frac{\partial^2 v_i}{\partial y^2} + \gamma/n \geq 0, \forall i,$$

which holds if and only if $\gamma \in [-\min_{i \in N}\left\{\frac{\partial^2 v_i}{\partial y^2}\right\}n, +\infty]$. ***Q.E.D.***

Therefore, when the punishment parameter is above the threshold, a large class of interesting learning dynamics converge, which is consistent with the experimental results. Intuitively, when the punishment parameter is sufficiently high, the incentive for each agent to match the mean of other agents' messages is also high. Therefore, when other agents increase their contributions, agent i also wants to increase her contribution to avoid the penalty. Thus the messages become strategic complements and the game is transformed into a supermodular game. Muench and Walker (1983) found a convergence condition for the Groves-Ledyard mechanism using Cournot best-reply dynamics and parameterized quadratic preferences. This proposition generalizes their result to general quasilinear preferences and a much wider class of learning dynamics.

Falkinger (1996) introduces a class of simple mechanisms. In this incentive compatible mechanism for public goods, Nash equilibrium is Pareto optimal when a parameter is chosen appropriately, i.e., when $\beta = 1 - 1/n$. However, it does not implement Lindahl allocations and the existence of equilibrium can be delicate in some environments.

DEFINITION 5. For the Falkinger (1996) mechanism, the strategy space of player i is $S_i \subset R^1$ with generic element $m_i \in S_i$. The outcome function of the public good and the net cost share of the private good for player i are

$$Y(m) = \sum_k m_k,$$

$$T_i^F(m) = b\left[m_i - \beta\left(m_i - \frac{\sum_{j \neq i} m_j}{n-1}\right)\right],$$

where $\beta > 0$.

This tax-subsidy scheme works as follows: if an individual's contribution is above the average contribution of the others, she gets a subsidy of β for a marginal increase in her contribution. If her contribution is below the average contribution of others, she has to pay a tax whereby a marginal increase in her contribution reduces her tax payment by β. If β is chosen appropriately, Nash equilibrium of this mechanism is Pareto efficient. Furthermore, it fully balances the budget both on and off the equilibrium path.

PROPOSITION 2. The Falkinger mechanism is a supermodular game for any $e \in E^{QD}$ if and only if $\beta \geq 1$.

Proof: Since u_i is C^2 on S_i, by Theorem 1, u_i has increasing differences in (m_i, m_{-i}) if and only if

$$\frac{\partial^2 u_i}{\partial m_i \partial m_j} = \frac{B_i b^2}{n-1}\beta(\beta - 1) \geq 0, \forall i,$$

which holds if and only if $\beta \geq 1$. Q.E.D.

Since Pareto efficiency requires that $\beta = 1 - 1/n$, in a large economy, this will produce a game which is close to being a supermodular game. It is interesting to note that in the quadratic environment of Falkinger et al. (2000), the game is very close to being a supermodular game: in the experiment β was set to 2/3. The results show the mean contribution level hovered around the Nash equilibrium, even though none of the 23 sessions had a mean contribution level exactly equal to the Nash equilibrium level in the last five rounds. Their results suggest that the convergence in supermodular games might be a function of the degree of strategic complementarity. That is, in games with a unique Nash equilibrium which can induce supermodular games, such as the Groves-Ledyard mechanism for any $e \in E^Q$ and the Falkinger mechanism for any $e \in E^{QD}$, as the degree of strategic complementarity increases, we might observe more rapid convergence to its stage game Nash equilibrium.

Three specific game forms implementing Lindahl allocations in Nash equilibrium have been introduced, Hurwicz (1979), Walker (1981), and Kim (1993). Since Tian (1989) and Peleg (1996) do not have specific mechanisms, we will only investigate the supermodularity of these three mechanisms. All three improve on the Groves-Ledyard mechanism in the sense that they all satisfy the individual rationality constraint in equilibrium. While Hurwicz (1979) and Walker (1981) can be shown to be unstable for any decentralized adjustment process in certain quadratic environments (Kim, 1986), the Kim mechanism is stable under a gradient adjustment process given quasilinear utility functions, which is a continuous time version of the Cournot-Nash tâtonnement adjustment process. Whether the Kim mechanism is stable under other decentralized learning processes is still an open question. Kim (1986) has shown that for any game form implementing Lindahl allocations there does not exist a decentralized adjustment process which ensures local stability of Nash equilibria in certain classes of environments.

PROPOSITION 3. None of the Hurwicz (1979), Walker (1981) and Kim (1993) mechanisms is a supermodular game for any $e \in E^Q$.

Proof: See Appendix. ∎

The following observation organizes all experimental results on Nash-efficient public goods mechanisms with available parameters by looking at whether they are supermodular games. The design parameters used in Smith's (1979) R1 treatment and Harstad and Marrese (1981) are not available.

OBSERVATION 1. *(1) None of the following experiments is a supermodular game: the Groves-Ledyard mechanism studied in Smith's (1979) R2 treatment, Harstad and Marrese (1982), Mori (1989), Chen and Plott (1996)'s low γ treatment, and Chen and Tang (1998)'s low γ treatment, the Walker mechanism in Chen and Tang (1998), and the Falkinger mechanism in Falkinger et al. (2000).*

(2) The Groves-Ledyard mechanism under the high γ in Chen and Plott (1996) and Chen and Tang (1998) are both supermodular games.

Therefore, none of the existing experiments which did not converge is a supermodular game, while those which did converge well are both supermodular games.

Note that designing a mechanism as a supermodular game might require some information on the part of the planner. For example, under the Groves-Ledyard mechanism, when choosing parameters to induce supermodularity, the planner needs to know the smallest second partial derivative of the players' utility for public goods in the society, i.e., $\min_{i \in N} \dfrac{\partial^2 v_i}{\partial y^2}$, for all possible levels of the public good, y, which is state-dependent information. In Nash implementation theory we usually assume that the planner does not have any information about the players' preferences. In that case, even though there exist a set of stable mechanisms among a family of mechanisms, the planner does not have the information to choose the right one. Therefore, in order to choose parameters to implement the stable set of mechanisms, the planner needs to have some information about the distribution of preferences and an estimate about the possible range of public goods level. One possible way of obtaining the information is through sampling (Gary-Bobo and Jaaidane, 2000). If the requisite information is not available, then an alternative might be to use "approximately" supermodular mechanisms, such as the Falkinger mechanism. In large economies when the planner selects $\beta = 1 - 1/n$ to induce efficiency, the mechanism is approximately supermodular.

6. CONCLUDING REMARKS

So far Nash implementation theory has mainly focused on establishing static properties of the equilibria. However, experimental evidence suggests that the fundamental question concerning any actual implementation of a specific mechanism is whether decentralized dynamic learning processes will actually converge to one of the equilibria promised by theory. Based on its attractive theoretical properties and the supporting evidence for these properties in the experimental literature, we focus on supermodularity as a robust stability criterion for Nash-efficient public goods mechanisms with a unique Nash equilibrium.

This paper demonstrates that given a quasilinear utility function the Groves-Ledyard mechanism is a supermodular game if and only if the punishment parameter is above a certain threshold while none of the Hurwicz, Walker and Kim mechanisms is a supermodular game. The Falkinger mechanism can be converted into a supermodular game in a quadratic environment if the subsidy coefficient is at least one. These results generalize a previous convergence result on the Groves-Ledyard mechanism by Muench and Walker (1983). They are consistent with the experimental findings of in Smith (1979), Harstad and Marrese (1982), Mori (1989), Chen and Plott (1996), Chen and Tang (1998), and Falkinger et al. (2000).

Two aspects of the convergence and stability analysis in this paper warrant attention. First, supermodularity is sufficient but not necessary for convergence to hold. It is possible that a mechanism could fail supermodularity but still behaves well on a class of adjustment dynamics, such as the Kim mechanism. Secondly, The stability analysis in this paper, like other theoretical studies of the dynamic stability of Nash mechanisms, have been mostly restricted to quasilinear utility functions. It is desirable to extend the analysis to other more general environments. The maximal domain of stable environments remains an open question.

Results in this paper suggest a new research agenda that systematically investigates the role of supermodularity in learning and convergence to Nash equilibrium. Two studies pioneer this new research agenda. Arifovic and Ledyard (2003) study the Groves-Ledyard mechanism in the same environment as Chen and Tang (1998), but use a much larger number of punishment parameters. Chen and Gazzale (forthcoming) study learning and convergence in Varian's (1994) compensation mechanism by systematically varying a free parameter below, close to, at and beyond the threshold of supermodularity to assess its effects on convergence. Findings from both studies are consistent. First, supermodular and "near-supermodular" games converge significantly better than those far below the threshold. Second, from a little below the threshold to the threshold, the improvement is statistically insignificant. Third, within the class of supermodular games, increasing the parameter far beyond the threshold does not significantly improve convergence. The robustness of these findings should be further investigated in future experiments in other games, for example, the Falkinger mechanism, as well as games outside the public goods domain.

ACKNOWLEDGMENT

I thank John Ledyard, David Roth and Tatsuyoshi Saijo for discussions that lead to this project; Klaus Abbink, Beth Allen, Rachel Croson, Roger Gordon, Elisabeth Hoffman, Matthew Jackson, Wolfgang Lorenzon, Laura Razzolini, Sara Solnick, Tayfun Sönmez, William Thomson, Lise Vesterlund, Xavier Vives, seminar participants in Bonn, Hamburg, Michigan, Minnesota, Pittsburgh, Purdue, and participants of the 1997 North America Econometric Society Summer Meetings (Pasadena, CA), the 1997 Economic Science Association meetings (Tucson, AZ), the 1998 Midwest

Economic Theory meetings (Ann Arbor, MI) and the 1999 NBER Decentralization Conference (New York, NY) for their comments and suggestions. The hospitality of the Wirtschaftspolitische Abteilung at the University of Bonn, the research support provided by Deutsche Forschungsgemeinschaft through SFB303 at the University of Bonn and NSF grant SBR-9805586 are gratefully acknowledged. Any remaining errors are my own.

NOTES

[1] A Lindahl equilibrium for the public goods economy is characterized by a set of personalized prices and an allocation such that utility and profit maximization and feasibility conditions are satisfied. As each consumer's consumption of the public good is a distinct commodity with its own market, externalities are eliminated. Thus, a Lindahl equilibrium is Pareto efficient. See, e.g., Milleron (1972).

[2] Note that the adaptive learning defined by Milgrom and Roberts (1990) does not include the simple reinforcement learning model of Roth and Erev (1995). It includes a subset of the EWA learning models (Camerer and Ho, 1999) for certain parameter combinations.

[3] "Theoretically, convergence implies that no deviation will ever be observed once the system equilibrates. In an experimental setting with long iterations, even after the system equilibrates, subjects sometimes experiment by occasional deviation. Therefore, it is necessary to have some behavioral definition of convergence: a system converges to an equilibrium at round t, if $x_i(s) = x_i^e$, $\forall i$ and $\forall s \geq t$, except for a maximum of n rounds of deviation for $s > t$, where n is small. For our experiments of 100 rounds, we let $n \leq 5$, i.e., there could be a total of up to 5 rounds of experimentation or mistakes after the system converged." (Chen and Tang, 1998).

REFERENCES

Arifovic, J. and Ledyard, J. (2003). "Computer Testbeds and Mechanism Design: Application to the Class of Groves-Ledyard Mechanisms for Provision of Public Goods." Manuscript, Caltech.

Boylan, R. and El-Gamal, M. (1993). "Fictitious Play: A Statistical Study of Multiple Economic Experiments." *Games Econ. Behavior* 5, 205–222.

Cabrales, A. (1999). "Adaptive Dynamics and the Implementation Problem with Complete Information." *Journal of Economic Theory* 86, 159–184.

Camerer, C. and Ho, T. (1999). "*Experienced-Weighted Attraction Learning in Normal Form Games*," *Econometrica*, Vol. 67, No. 4. (Jul., 1999), pp. 827–874.

Chen, Y. (forthcoming). "Incentive-Compatible Mechanisms for Pure Public Goods: A Survey of Experimental Research," in *The Handbook of Experimental Economics Results* (C. Plott and V. Smith, Eds.). Amsterdam: Elsevier Press.

Chen, Y. and Gazzale, R. (forthcoming). "When Does Learning in Games Generate Convergence to Nash Equilibria? The Role of Supermodularity in an Experimental Setting." *American Economic Review*.

Chen, Y. and Plott, C. (1996). "The Groves-Ledyard Mechanism: An Experimental Study of Institutional Design." *Journal of Public Economics* 59, 335–364.

Chen, Y. and Tang, F. (1998). "Learning and Incentive-Compatible Mechanisms for Public Goods Provision: An Experimental Study." *Journal of Political Economy* 106, 633–662.

Dibvig, P. and Spatt, C. (1983). "Adoption Externalities as Public Goods." *Journal of Public Economics* 20, 231–247.

El-Gamal, M. and Grether, D. (1995). "Uncovering Behavioral Strategies: Are People Bayesians?" *Journal of the American Statistical Associations* 90, 1137–1145.

Falkinger, J. (1996). "Efficient Private Provision of Public Goods by Rewarding Deviations from Average." *Journal of Public Economics* 62, 413–422.

Falkinger, J., Fehr, E., Gächter, S. and Winter-Ebmer, R. (2000). "A Simple Mechanism for the Efficient Provision of Public Goods – Experimental Evidence." *American Economic Review* 90, 247–264.

Gary-Bobo, R. and Jaaidane, T. (2000). "Polling Mechanisms and the Demand Revelation Problem." *Journal of Public Economics* 76, 203–238.

Green, J. and Laffont, J.-J. (1977). "Characterization of Satisfactory Mechanisms for the Revelation of the Preferences for Public Goods." *Econometrica* 45, 427–438.

Groves, T. and Ledyard, J. (1977). "Optimal Allocation of Public Goods: A Solution to the 'Free Rider' Problem." *Econometrica* 45, 783–809.

Harstad, R. and Marrese, M. (1981). "Implementation of Mechanism by Processes: Public Good Allocation Experiments." *Journal of Economic Behavior and Organization* 2, 129–151.

Harstad, R. and Marrese, M. (1982). "Behavioral Explanations of Efficient Public Good Allocations." *Journal of Public Economics* 19, 367–383.

Hurwicz, L. (1972). "On Informationally Decentralized Systems," in *Decision and Organization* (C. McGuire and R. Radner, Eds.), pp. 297–336. Amsterdam: North Holland Press.

Hurwicz, L. (1979). "Outcome Functions Yielding Walrasian and Lindahl Allocations at Nash Equilibrium Points." *Review of Economic Studies* 46, 217–225.

Kim, T. (1986). "On the Nonexistence of a Stable Nash Mechanism implementing Lindahl Allocations." Manuscript: University of Minnesota.

Kim, T. (1993). "A Stable Nash Mechanism Implementing Lindahl Allocations for Quasi-linear Environments." *Journal of Mathematical Economics* 22, 359–371.

Milgrom, P. and Roberts, J. (1990). "Rationalizability, Learning and Equilibrium in Games with Strategic Complementarities." *Econometrica* 58, 1255–1277.

Milgrom, P. and Roberts, J. (1991). "Adaptive and Sophisticated Learning in Normal Form Games." *Games Econ. Behavior* 3, 82–100.

Milleron, Jean-Claude (1972). "Theory of Value with Public Goods: A Survey Article." *Journal of Economic Theory* 5, 419–477.

Mori, T. (1989). "Effectiveness of Mechanisms for Public Goods Provision: An Experimental Study." *Economic Studies* 40, 234–246.

Muench, T. and Walker, M. (1983). "Are Groves-Ledyard Equilibria Attainable?" *Review of Economic Studies* **L**, 393–396.

Moulin, H. (1984). "Dominance Solvability and Cournot Stability." *Mathematical Social Sciences* 7, 83–102.

Peleg, B. (1996). "Double Implementation of the Lindahl Equilibrium by a Continuous Mechanism." *Economic Design* 2, 311–324.

Roberts, J. (1979). "Incentives and Planning Procedures for the Provision of Public Goods." *Review of Economic Studies* 46, 283–292.

Roth, A. and Erev, I. (1995). "Learning in Extensive Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." Games and Economic Behavior 8: 164–212.

Smith, V. (1979). "Incentive Compatible Experimental Processes For the Provision of Public Goods," in *Research in Experimental Economics* 1 (V. Smith, Eds.), pp. 59–168. Greenwich, CT: JAI Press.

Tian, G. (1989). "Implementation of the Lindahl Correspondence by a Single-Valued, Feasible, and Continuous Mechanism." *Review of Economic Studies* 56, 613–621.

Topkis, D. (1978). "Minimizing a Submodular Function on a Lattice." *Operations Research* 26, 305–321.

Topkis, D. (1979). "Equilibrium Points in Nonzero-Sum n-Person Submodular Games." *SIAM Journal of Control and Optimization* 17, 773–787.

de Trenqualye, P. (1988). "Stability of the Groves-Ledyard Mechanism." *Journal of Economic Theory* 46, 164–171.

de Trenqualye, P. (1989). "Stable Implementation of Lindahl Allocations." *Economic Letters* 29, 291–294.

Vega-Redondo, F. (1989). "Implementation of Lindahl Equilibrium: An Integration of Static and Dynamic Approaches." *Mathematical Social Sciences* 18, 211–228.

Walker, M. (1980). "On the Impossibility of a Dominant Strategy Mechanism to Optimally Decide Public Questions." *Econometrica* 48, 1521–1540.

Walker, M. (1981). "A Simple Incentive Compatible Scheme for Attaining Lindahl Allocations." *Econometrica* 49, 65–71.

APPENDIX

Before proving Proposition 3, we first define the three mechanisms. All three mechanisms require that the number of players be at least three, i.e., $n \geq 3$.

DEFINITION 6. *For the Hurwicz (1979) mechanism, the strategy space of player i is $S_i \subset R^2$ with generic element $(p_i, y_i) \in S_i$. The outcome function of the public good and the net cost share of the private good for player i are*

$$Y(y) = \frac{\sum_k y_k}{n},$$

$$T_i^H(p, y) = R_i \cdot Y(y) + p_i \cdot (y_i - y_{i+1}) - p_{i+1}(y_{i+1} - y_{i+2})^2,$$

where $R_i = \frac{1}{n} + p_{i+1} - p_{i+2}$.

DEFINITION 7. *For the Walker (1981) mechanism, the strategy space of player i is $S_i \subset R^1$ with generic element $m_i \in S_i$. The outcome function of the public good and the net cost share of the private good for player i are*

$$Y(m) = \sum_k m_k,$$

$$T_i^W(m) = \left(\frac{1}{n} + m_{i-1} - m_{i+1}\right) \cdot Y(m).$$

DEFINITION 8. *For the Kim (1993) mechanism, the strategy space of player i is $S_i \subset R^2$ with generic element $(m_i, z_i) \in S_i$. The outcome function of the public good and the net cost share of the private good for player i are*

$$Y(m) = \sum_k m_k,$$

$$T_i^K(m, z) = P_i(m, z) \cdot Y(m) + \frac{1}{2}\left(z_i - \sum_k m_k\right)^2,$$

where $P_i(m, z) = \frac{b}{n} - \sum_{j \neq i} m_j + \frac{1}{n} \sum_{j \neq i} z_j$.

Proof of Proposition 3:

(1) To show that the Hurwicz mechanism is not a supermodular game for any $e \in E^Q$, it suffices to show that the payoff function, u_i, does not have increasing difference in (y_i, y_{-i}).

Since $u_i(p, y) = v_i(y) + \omega_i - T_i^H$, we have

$$\frac{\partial^2 u_i}{\partial y_i \partial y_j} = \frac{1}{n^2}\frac{\partial^2 v_i}{\partial y^2}, \quad \text{for all } j \neq i + 1.$$

By Definition 1, $\frac{\partial^2 v_i}{\partial y^2} < 0$, so

$$\frac{\partial^2 u_i}{\partial y_i \partial y_j} < 0, \quad \text{for all } i \text{ and for all } j \neq i + 1.$$

By Theorem 1, u_i does not have increasing difference in (y_i, y_{-i}).

(2) To show that the Walker mechanism is not a supermodular game for any $e \in E^Q$, it suffices to show that the payoff function, u_i, does not have increasing difference in (m_i, m_{-i}):

$$\frac{\partial^2 u_i}{\partial m_i \partial m_j} = \begin{cases} \frac{\partial^2 v_i}{\partial y^2} + 1, & \text{if } j = i + 1; \\ \frac{\partial^2 v_i}{\partial y^2} - 1, & \text{if } j = i - 1; \\ \frac{\partial^2 v_i}{\partial y^2}, & \text{if } j \neq i - 1, i + 1. \end{cases}$$

By Definition 1, $\frac{\partial^2 v_i}{\partial y^2} < 0$, so

$$\frac{\partial^2 u_i}{\partial m_i \partial m_j} < 0 \quad \text{for all } j \quad i + 1.$$

By Theorem 1, u_i does not have increasing difference in (m_i, m_{-i}).

(3) To show that the Kim mechanism is not a supermodular game for any $e \in E^Q$, it suffices to show that the payoff function, u_i, does not have increasing difference in (m_i, z_{-i}):

$$\frac{\partial^2 u_i}{\partial m_i \partial z_j} = -\frac{1}{n} < 0.$$

By Theorem 1, u_i, does not have increasing difference in (m_i, z_{-i}). ***Q.E.D.***

Chapter 11

ENTRY TIMES IN QUEUES WITH ENDOGENOUS ARRIVALS: DYNAMICS OF PLAY ON THE INDIVIDUAL AND AGGREGATE LEVELS

J. Neil Bearden
University of Arizona

Amnon Rapoport
University of Arizona
and
Hong Kong University of Science and Technology

Darryl A. Seale
University of Nevada Las Vegas

Abstract

This chapter considers arrival time and staying out decisions in several variants of a queueing game characterized by endogenously determined arrival times, simultaneous play, finite populations of symmetric players, discrete strategy spaces, and fixed starting and closing times of the service facility. Experimental results show 1) consistent patterns of behavior on the aggregate level in all the conditions that are accounted for quite well by the symmetric mixed-strategy equilibrium of the stage game, 2) considerable individual differences in arrival time distributions that defy classification, and 3) learning trends across iterations of the stage queueing game in some of the experimental conditions. We propose and subsequently test a simple reinforcement-based learning model that, with a few exceptions, accounts for these major findings.

Keywords: Queueing, Endogenous Arrivals, Equilibrium Analysis, Experimentation
JEL Classification: C72, C92

1. INTRODUCTION

In two recent experiments, Rapoport, Stein, Parco, and Seale (RSPS, in press) and Seale, Parco, Stein, and Rapoport (SPSR, 2003) have studied arrival time and

A. Rapoport and R. Zwick (eds.), Experimental Business Research, Vol. II, 201–221.

staying out decisions in a class of queueing problems with endogenously determined arrival times, a finite and commonly known calling population of players ($n = 20$ in both experiments), discrete strategy spaces, and fixed starting and closing time of the service facility. Focusing on transient behavior, these problems differ from the ones typically studied in queueing theory that assume continuous strategy spaces, steady-state behavior, and exogenously determined arrival times (but see Hassin & Haviv, 2003; Lariviere & Mieghem, 2003). The objective of each player in the queueing problems studied by RSPS and SPSR is to maximize her expected payoff by completing service while minimizing her waiting time in the queue. Formulating these queueing problems as complete information, non-cooperative games in strategic form, RSPS and subsequently SPSR constructed a Markov chain algorithm to compute symmetric mixed-strategy equilibrium solutions to the queueing games. Implementing a repeated game design, they then assessed the descriptive power of these solutions in several variants of the game. These variants differ from one another on three dimensions: 1) whether arrivals before the starting time of the service facility are allowed; 2) whether all the n players can complete their service with no waiting in line; and 3) whether at the end of each stage game (trial) players only receive private information about their own outcome or public information about the decisions and payoffs of all the n players.

Using several statistics to compare observed to equilibrium (predicted) behavior (e.g., mean payoffs, distribution of arrival times, distribution of interarrival times, distribution of waiting times in the queue), RSPS and SPSR reported three major findings. First, with one exception that we discuss later, they reported consistent patterns of behavior on the aggregate level that can be accounted for remarkably well by the symmetric mixed-strategy equilibrium. Second, they observed considerable individual differences in arrival time and staying out decisions that defied classification. Most subjects often switched their decisions from trial to trial but definitely not in a manner consistent with equilibrium play. Third, they reported learning trends across trials that strongly depended on the dimensions mentioned above. In particular, when the parameter values of the game were so selected that all the n players could, in principle, complete service without waiting in line (and, consequently, maximize their individual payoffs), there was only very weak evidence for learning across trials regardless of the nature of the outcome feedback (private vs. group) that was provided at the end of each trial. When the parameter values were selected so that in equilibrium a substantial fraction of the players should stay out on each trial, learning depended on the nature of the outcome feedback. If each player was informed at the end of each trial of the decisions (staying out or arrival time) and payoffs of all the n players, then SPSR reported strong evidence of learning in the direction of equilibrium play with most players first receiving negative payoffs because of congestion (not enough players staying out) and then gradually approaching equilibrium play by increasing the frequency of staying out decisions. If each player was only informed of his own decision and payoff, learning did not take place and most of the players ended up deep in the negative payoff domain.

The major purpose of the present paper is to explain and reconcile these three major findings. Focusing on the dynamics of play, we present and then test a simple model in an attempt to explain 1) how the aggregation of individual arrival time distributions that differ considerably from one another results in replicable patterns that are accounted for by the symmetric mixed-strategy equilibrium, and 2) how outcome information (private vs. public) affects learning when the service facility cannot accommodate all the members of the calling population between its starting and closing times. Although the analyses that we present below mostly focus on the individual and aggregate distributions of arrival time (that also include the decision to stay out of the queue), we also comment on the distribution of frequency of switching the decision from one trial to another and the distribution of the magnitude of such switches.

The rest of the chapter is organized as follows. Section 2 states the assumptions of the queueing game and illustrates it with an example. Section 3 describes the mixed-strategy equilibrium distributions of arrival time for the three variants of the queueing game studied by RSPS and SPSR, and then compares them to observed aggregate distributions. Selected individual distributions of arrival time are also presented both to illustrate the differences among members of the same population and the failure of the mixed-strategy equilibrium to account for individual behavior. Section 4 describes a simple reinforcement-type learning model and the estimation of its parameters. Section 5 contains a discussion of the model's success or failure in accounting for the three major findings mentioned above. Section 6 concludes.

2. THE QUEUEING GAME WITH ENDOGENOUS ARRIVAL TIME

The queueing game is characterized by a 6-tuple (n, d, c, r, g, T), where n is the number of players and d is the (fixed) time required to serve a single player (same for all n players). There are three payoff parameters, namely, c, r, and g: c is the per unit waiting cost, r is the payoff for completing service, and g is the payoff for staying out of the queue. $T + 1$ is the number of entry periods (pure strategies). For example, if the service facility is open for exactly two hours and entry time is measured in minutes, then there are $T + 1 = 121$ possible entry times.

The following assumptions characterize the game. The service facility opens at T_o and closes at T_e. Arrivals are made in discrete time units (single minutes is RSPS and 5 minute intervals in SPSR). Decisions are made simultaneously and anonymously. Thus, at the beginning of each trial, each player must decide whether to enter the queue. If she decides to do so, she must specify her time of arrival (e.g., 8:01, 8:02, . . . in RSPS). If m players happen to arrive at the same time, $2 \le m \le n$, then their order of arrival is determined randomly with equal probability $1/m$. Balking (not entering the queue upon arrival) and reneging (departing the queue after arrival and before service commences) are prohibited. One implication of the latter rule is that players cannot leave the queue even if they know with certainty that service will not be provided. Early arrivals before time T_o may (SPSR) or may not

(RSPS) be allowed. Service time for each player, d, is fixed, and the queue discipline is FIFO. There is a single server, a single service stage, and no limit on the queue length. Because the decisions are made simultaneously, players cannot observe the state of the queue before making their decisions. Finally, the payoff function – the same for all n players – is given by

$$H_i = \begin{cases} g & \text{if player } i \text{ stays out} \\ -c \times w_i & \text{if player } i \text{ waits } w_i \text{ times units and fails to complete service} \\ r - c \times w_i & \text{if player } i \text{ waits } w_i \text{ times units and complets service} \end{cases}$$

where w_i is the time spent in the queue until service commences. No waiting cost is charged for the time (d) spent being served. RSPS and SPSR make the natural assumptions: $r > g$, $r > c$, and $c > 0$. The values of T_o, T_e and d, as well as the values of the waiting times w_i, are measured in commensurate units. The three payoff parameters c, r, and g, the population size n, and the opening and closing times T_o and T_e are assumed to be commonly known. For a discussion of the assumptions and their justification see RSPS and SPSR.

Example Table 1 provides an example that illustrates the queueing game and the computation of the individual payoffs. (See the subject instructions in the Appendix of RSPS for a similar example.) The parameter values for this example are $n = 20$, $d = 45$, $T_o = 8{:}00$, $T_e = 18{:}00$, $c = 1$, $r = 100$, and $g = 15$. The same parameter values are used in two of the four conditions in SPSR (see below). Players are restricted to arrive at 5-minute time intervals, and early arrivals (before T_o) are allowed. Payoffs are in pennies.

Columns 1 and 2 of Table 1 present the player number (an integer from 1 to 20) and the players' decisions. In this example, 16 of the 20 players opted to enter (at possibly different times), and 4 players (6, 16, 11 and 4, who are listed at the bottom) decided to stay out. Players 13, 3, and 18 arrived at 7:05, 7:25, and 7:30, before the opening time T_o, and had to wait in line 55, 80, and 120 minutes, respectively. All three completed service. Players 15 and 1 arrived at exactly 8:00, and the two-player tie was resolved in favor of player 15 (who still had to wait $45 \times 3 = 135$ minutes until players 13, 3, and 18 completed service). Player 14 arrived at 14:55 and was served immediately with no delay. Although players 2, 17, and 7 arrived more than 45 minutes before closing time, none of them received service. Of the twenty players in this example, eight lost money. Total system idle time was 25 minutes, from 14:45 to 14:55 and from 17:45 to 18:00. Columns 3, 4, and 5 present the beginning of the service time, the waiting time (in minutes), and the waiting costs. The reward (that could assume one of the three values r, 0, or g) is presented in column 6, and the payoff is shown in the right-hand column.

Table 1. Example of a Queueing Game when Early Arrivals are Possible (d = 45)

Player	*Decision*	*Service Starts at*	*Waiting Time*	*Waiting Cost*	*Reward*	*Payoff*
13	Arrive: 7:05	8:00	55	$0.55	$1.00	$0.45
3	Arrive: 7:25	8:45	80	0.80	1.00	0.20
18	Arrive: 7:30	9:30	120	1.20	1.00	−0.20
15	Arrive: 8:00	10:15	135	1.35	1.00	−0.35
1	Arrive: 8:00	11:00	180	1.80	1.00	−0.80
5	Arrive: 8:45	11:45	180	1.80	1.00	−0.80
20	Arrive: 10:00	12:30	150	1.50	1.00	−0.50
10	Arrive: 12:10	13:15	65	0.65	1.00	0.35
8	Arrive: 13:45	14:00	15	0.15	1.00	0.85
14	Arrive: 14:55	14:55	0	0	1.00	1.00
9	Arrive: 15:00	1:40	40	0.40	1.00	0.60
12	Arrive: 15:00	16:25	85	0.85	1.00	0.15
19	Arrive: 15:30	17:10	100	1.00	1.00	0
2	Arrive: 16:25	NA	100	1.00	0	−1.00
17	Arrive: 17:00	NA	85	0.85	0	−0.85
7	Arrive: 17:00	NA	60	0.60	0	−0.60
6	Stay out	None	0	0	0.15	0.15
16	Stay out	None	0	0	0.15	0.15
11	Stay out	None	0	0	0.15	0.15
4	Stay out	None	0	0	0.15	0.15

3. PREDICTED AND OBSERVED RESULTS

3.1. Experimental Conditions

RSPS and SPSR together conducted three different experimental conditions that differ from one another in one or more parameters or assumptions. These conditions are described below. In all three conditions $n = 20$ and the number of iterations of the stage game is 75. All the experiments are computer-controlled.

Condition 1 (RSPS). $T_o = 8{:}00$, $T_e = 18{:}00$, $d = 30$, $c = 1$, $r = 60$, and $g = 0$. Time is measured in single minute intervals, and early arrivals are prohibited. This parameterization gives rise to 601 possible entry times, namely 8:00, 8:01, . . . , 18:00, and another decision of staying out. Information is private. In particular, at the end of each trial each player is reminnded of her decision (arrival time or staying out); number of players tied at her arrival time, if any; and the outcome of the tie-breaking rule; her queue waiting time (w_i); her payoff for the *trial* (H_i); and her cumulative payoff from the beginning of the session. We refer to this information condition as *Private Outcome Information.*

Condition 2 (SPSR). $T_o = 8{:}00$, $T_e = 18{:}00$, $d = 30$, $c = 1$, $r = 100$, $g = 15$. Time is measured in 5-minute intervals, and early arrivals are allowed. To limit the strategy space, players are not allowed to enter the queue before 6:00. In fact, this requirement imposes no practical limitation. This parameterization gives rise to 145 possible entry times, namely, 6:00, 6:05, . . . , 18:00 and an additional pure strategy of staying out.

Condition 2 was further divided into two sub-conditions, namely Condition 2P and Condition 2G, in terms of the information provided to the player at the end of each trial. Condition 2P included *Private Outcome Information.* Condition 2G included *Group Outcome Information* which consisted, in addition to the *Private Outcome Information*, of complete information about the 1) arrival times and staying out decisions, 2) service starting time, and 3) individual payoffs for all the n players in the session. This was accomplished by presenting a computer "Results" screen at the end of each trial that consisted of a 20×3 matrix with rows corresponding to the twenty players arranged in terms of the time of their arrival (staying out decisions were placed at the bottom rows), and three columns corresponding to the player's decision (arrival time or staying out), starting time of service, and individual payoff for the trial (see Appendix in SPSR for details).

Condition 3 (SPSR). Condition 3 used the same parameter values as Condition 2 with the only exception that $d = 45$ minutes. It, too, was further divided into two sub-conditions, Condition 3P and Condition 3G that incorporated the Private and Group Outcome Information, respectively. Note that if $d = 30$ (Conditions 1 and 2), all the n players can complete service without waiting if they arrive at 30 minute intervals starting exactly at 8:00. In contrast, only 13 of the 20 players can complete service in Condition 3 without waiting in the queue, if they arrive at 45 minute intervals starting at exactly 8:00 (8:00, 8:45, . . . , 17:00), whereas the remaining 7 players have to stay out. As we show below, this difference in service time and whether or not early arrivals are allowed strongly affect the mixed-strategy equilibria for these three experimental conditions.

3.2. Method

Subjects. Condition 1 included four groups of $n = 20$ members each, whereas Conditions 2P, 2G, 3P, and 3G each included two groups of $n = 20$ players for a total of 12 groups (240 subjects) across conditions. With the exception of Group 4 in Condition 1, all the subjects were University of Arizona students, mostly undergraduates, who volunteered to participate in a decision making experiment for payoff contingent on performance. Males and females participated in almost equal proportions. Group 4 in Condition 1 included twenty "sophisticated" subjects who participated in a summer workshop on experimental economics that was conducted at the University of Arizona. Members of this group were graduate students and post-doctoral fellows of economics with a keen interest in experimental economics and solid background in game theory. Individual payoff ranged considerably, depending on the experimental condition, from $15.00 to $53.24. The conversion rate of the fictitious currency (called "francs") used in the experiment was doubled for the "sophisticated" subjects in Group 4 of Condition 1.

Procedure. Details of the experimental procedure appear in RSPS and SPSR and will not be repeated here. Basically, in all three conditions the queueing game was presented as an emissions control scenario with a fixed and commonly known number of car owners, a station whose opening and closing times are fixed and commonly known, fixed service time per customer, and a common payoff structure (see above). At the beginning of the session, each subject was provided with an endowment of 1,000 francs. Francs earned during each trial were added or subtracted from this endowment. At the end of the session, the cumulative payoff in francs was converted to US dollars (100 francs = US$1.00). Subjects who ended the session losing their entire endowment were only paid their show-up fee. Subjects were paid individually and dismissed.

Equilibrium Analysis. Each of the queueing games in Conditions 1 and 2 has $n!$ equilibria in pure strategies, where players arrive at 30 minute intervals starting at 8:00. Under pure strategy equilibrium play, each player has zero waiting time with an associated payoff of r. Technically, these are coordination games with $n!$ pure-strategy equilibria that are not Pareto-rankable and do not depend on the reward to cost ratio r/c. Without pre-play communication, coordination on any one pure strategy equilibrium is practically impossible due the large size of the group even under multiple iterations of the stage game. The queueing game in Condition 3 has multiple asymmetric pure-strategy equilibria where 13 players enter the queue (with at least 45 minute intervals between consecutive arrivals and, consequently, no waiting time) and 7 players staying out. Again, it is highly unlikely that the twenty players could coordinate on any particular equilibrium, even in Condition 3G, without pre-play communication.

Each of the three queueing games in Conditions 1, 2, and 3 has a symmetric mixed-strategy equilibrium solution. The Appendix of RSPS contains a detailed description of the computational method used to construct these solutions. Essentially, it consists of specifying the state space, the transitional probabilities of the stochastic process that governs the queue progression, and the iterative procedure to compute the arrival times and staying out decisions under mixed-strategy play. Figs. 1, 2, and 3 display the equilibrium solutions for the three games in Conditions 1, 2, and 3, respectively.

Several features of the equilibrium solutions warrant discussion. In the solution for Condition 1 (Fig. 1), players join the queue at the earliest possible time of 8:00 ($t = 0$) with probability 0.211 and stay out with probability 0.060. They should never join the queue between 8:01 and 9:03, and then join the queue with positive probabilities until 17:30 ($t = 570$). Because $g = 0$ in Condition 1, the expected payoff under this equilibrium is clearly zero. In the equilibrium solution for Condition 2 (Fig. 2), players should always enter the queue starting at 6:35 and ending at 17:30. The expected payoff under equilibrium play is $18.35 > g = 15$. In contrast, the equilibrium solution for Condition 3 displays a very different pattern. The probability of staying out is 0.4096, implying that, on average, 8.2 players out of 20 should stay out on each trial. The expected value is clearly $g = 15$.

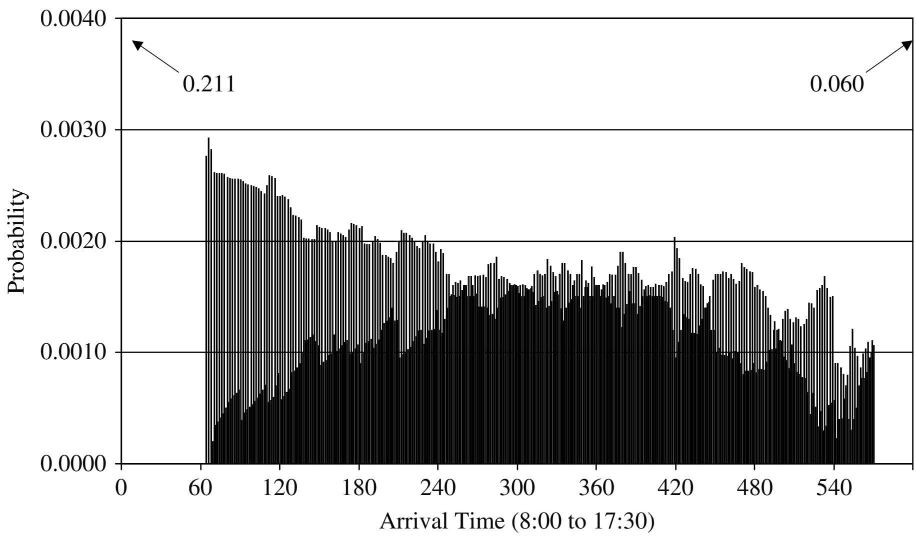

Figure 1. Symmetric mixed-strategy equilibrium of arrival times for Condition 1.

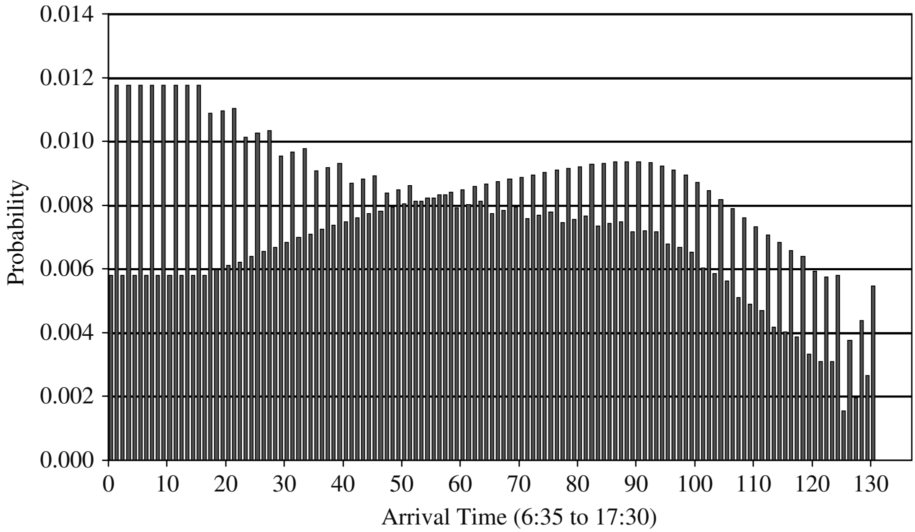

Figure 2. Symmetric mixed-strategy equilibrium of arrival times for Condition 2.

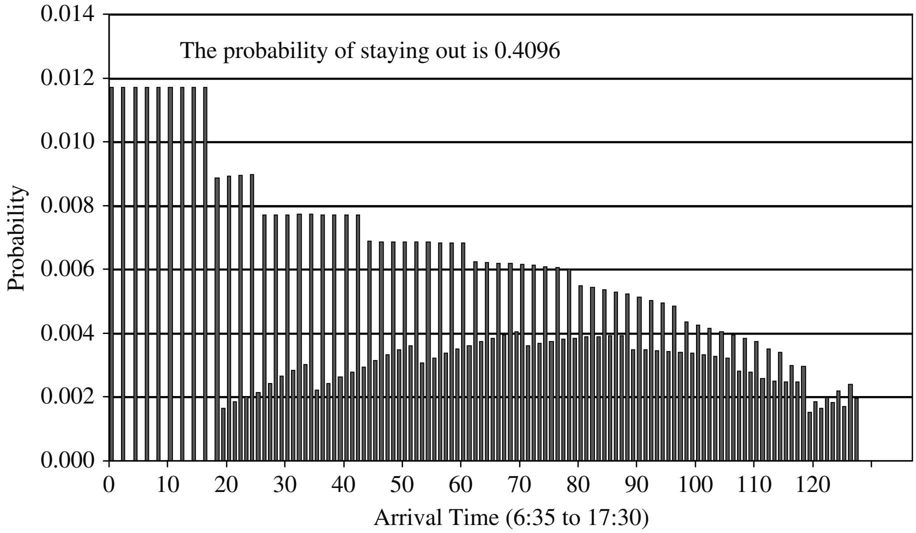

Figure 3. Symmetric mixed-strategy equilibrium of arrival times for Condition 3.

All three figures exhibit the periodicity of the solution that has also been reported in solving for the equilibria of games with smaller strategy spaces and smaller number of players. For example, in Fig. 2 players should arrive at the queue at 6:40, 6:50, . . . , 8:00 with probability 0.0058 and at the intermediate times 6:45, 6:55, . . . , 7:55 with probability 0.0118, which is twice as large. In Fig. 3, players should enter the queue at times 6:40, 6:50, . . . , 8:00 with probability 0.0117 and stay out at intermediate times 6:45, 6:55, . . . , 8:05. This periodicity is due to a combination of the discretization of the strategy space, fixed service time, and fixed opening (T_o) and closing times (T_e).

3.3. Results

Observed Arrival Time Distributions: Aggregate Results. Using several different statistics, RSPS reported no significant differences among the four groups in Condition 1. In particular, although the "sophisticated" subjects in Group 4 were paid twice as much as the other subjects (and took about twice as much time to complete the session), their results did not differ from those of the other three groups. Therefore, the results of all four groups were combined ($4 \times 20 \times 75 = 6000$ observations). Fig. 4 displays the observed and predicted (equilibrium) cumulative probability distributions of arrival time (staying out decisions are treated as arrivals at time 18:00). The statistical comparison of observed and predicted arrival time distributions is problematic because of the dependencies between and within players. Strictly speaking, the group is the unit of analysis, resulting in only four degrees of freedom for the statistical comparison. The one-sample two-tailed Kolmogorov-Smirnov (K-S) test (df = 4) could not reject the null hypothesis of no difference between

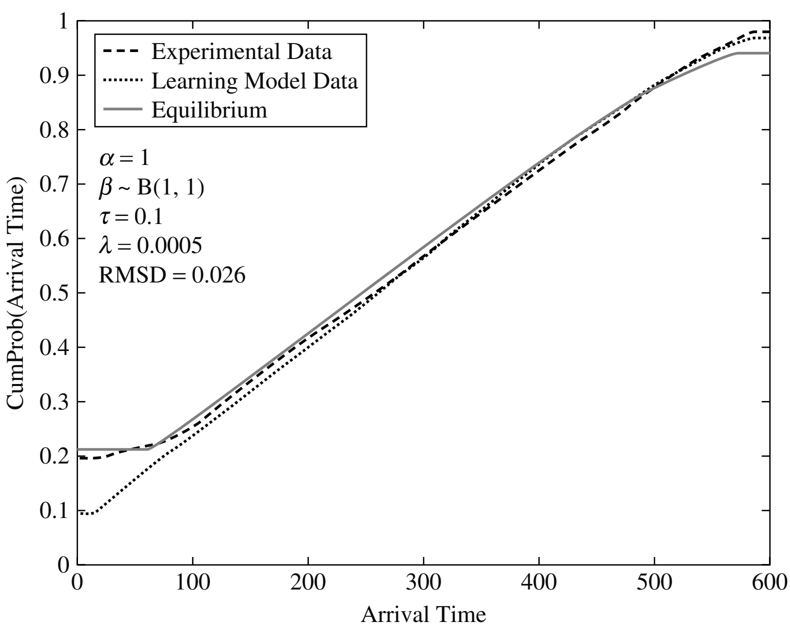

Figure 4. Observed and predicted distribution of arrival time and staying out decisions in Condition 1.

the observed and predicted distributions of arrival time. Assuming independence between (but not within) subjects yielded df = 80. But even with this considerably more conservative test, the same null hypothesis could not be rejected ($p > 0.05$). RSPS detected three minor discrepancies between observed and predicted probabilities of arrival time in all four groups (see Fig. 4): 1) the observed proportion of arriving at exactly 8:00 was smaller (by 0.02) than predicted; 2) the observed proportion of arriving between 8:01 and 9:03 was 0.031 compared to the theoretical value of zero; 3) the proportion of staying out was smaller than predicted. A more detailed analysis that broke the 75 trials into three blocks of 25 trials each shows that the first two discrepancies decreased across blocks in the direction of equilibrium play.

SPSR similarly reported no significant differences between the two groups in Condition 2G. Of the four tests used in this comparison, two yielded statistical differences between the two groups in Condition 2P. Nevertheless, the results were also combined across these two groups. Using the same format as Fig. 4, Fig. 5 exhibits the observed and predicted cumulative distributions of arrival time for Condition 2P (upper panel) and Condition 2G (lower panel). Similarly to Condition 1, the K-S test could not reject the null hypothesis of no difference between the observed and predicted distributions of arrival time (D = 0.059 for Condition 2G,

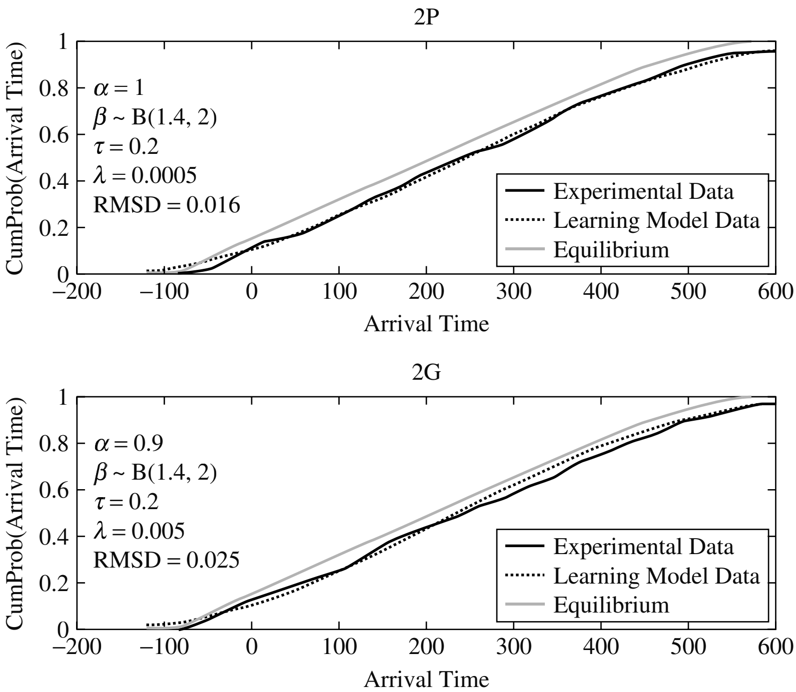

Figure 5. Observed and predicted distribution of arrival time and staying out decisions in Condition 2.

and $D = 0.069$ for Condition 2P; $n = 40$ and $p > 0.05$ in each case) even under the conservative assumption of independence between subjects. Notwithstanding these results, Fig. 5 shows two minor but systematic discrepancies between observed and predicted distributions of arrival time: 1) the observed proportion of entry before 7:35 was smaller than predicted; 2) approximately 4% of all the decisions were to stay out compared to 0% under equilibrium play. A more detailed analysis that breaks the 75 trials into three blocks shows that the former discrepancy decreased across trials but the latter did not. Analyses of individual data show that a few subjects stayed out on 6 or more (out of 75) trials either in an attempt to take time to consider their future decisions or to increase their cumulative payoff (by g) after a sequence of losses.

Turning next to Condition 3, SPSR also reported no significant differences between the two groups in Condition 3G and no significant differences between the two groups in Condition 3P. The two groups in each of these two conditions were separately combined to compute the aggregate distributions of arrival times. Fig. 6 portrays the observed and predicted cumulative distributions of arrival time for Condition 3P (upper panel) and 3G (lower panel). The K-S test once again could not reject the null hypothesis of no differences between the two distributions ($D = 0.061$

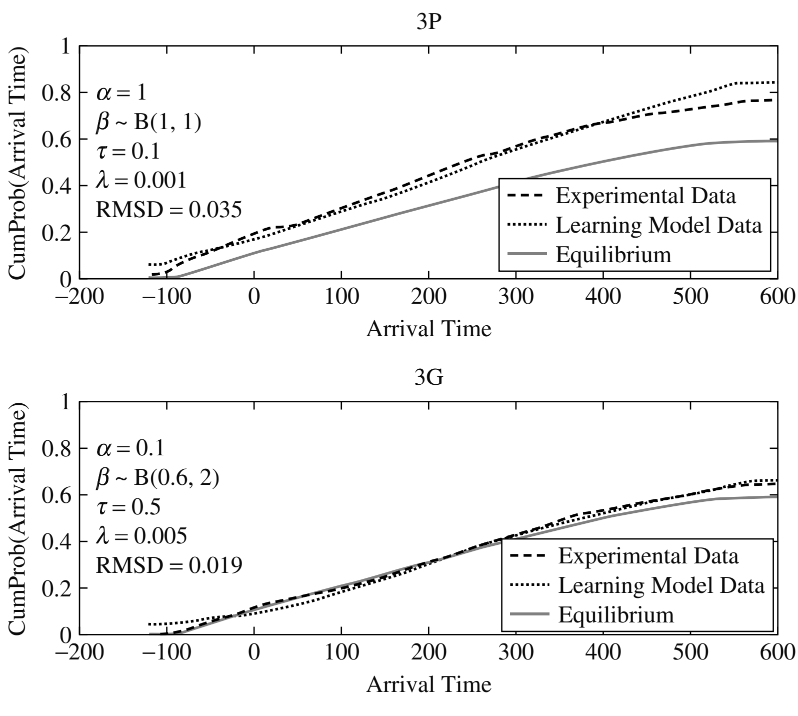

Figure 6. Observed and predicted distribution of arrival time and staying out decisions in Condition 3.

and $D = 0.121$ for Conditions 3G and 3P, respectively; $n = 40$ and $p > 0.05$ in each case). Nevertheless, the upper panel shows that subjects in Condition 3P did not stay out as frequently as predicted. A further analysis that focuses on the staying out decisions shows that the percentage of staying out decisions in Condition 3G steadily increased from 30% in trials 1–25 through 35.5% in trials 26–50 to 40.5% in trials 51–75. Compare the latter percentage to the equilibrium percentage of 40.96%. In contrast, there was no evidence for learning across blocks of trials in Condition 3P. As the subjects in Condition 3P received no information on the number of subjects staying out on any given trial, they had no way of determining whether their payoff for the trial – which was typically negative – was due to a poor choice of entry time or insufficient number of staying out decisions. This was not the case in Condition 3G, where Group Outcome Information was provided. Subjects in Condition 3G, who often lost money on the early trials, used this information to slowly recover their losses by having more (but not necessarily the same) subjects staying out on each trial. In contrast, most of the subjects in Condition 3P entered the queue more frequently than predicted and consequently almost never recovered their losses.

Observed Arrival Time Distributions: Individual Results. In contrast to the aggregate distributions of arrival time that show remarkable consistency across groups and are accounted for quite well by the equilibrium solution, the individual distributions of arrival time differ considerably from one another, show no support for mixed-strategy equilibrium play, and defy a simple classification. One representative group – Group 1 of Condition 1 – was selected to illustrate the contrast between the consistent patterns of arrival on the aggregate level and heterogeneous patterns of arrival on the individual level. Fig. 7 exhibits the individual arrival times of all the 20 subjects in Group 1 of Condition 1. We have opted to display the arrival times by trial rather than combine them into frequency distributions. Thus, the horizontal axis in each individual display counts the trial number from 1 through 75, and the vertical axis shows the arrival time on a scale from 6:00 (bottom) to 18:00 (top). A short vertical line that extends below the horizontal axis (i.e., below 0) indicates no entry. We observe that Subject 5 (first from left on row 2), after switching her entry time, entered at 8:00 on all trials after trial 25. In contrast, Subject 13 (first from left on row 4) never entered the queue at 8:00. Subject 9 (first from left on row 3) stayed out on 10 of the 75 trials, whereas Subjects 1, 2, 5, 6, 7, 8, 11, 13, 14, 17, and 18 never stayed out. Most of the staying out decisions is due to Subjects 9 and 15.

4. QUEUING LEARNING MODEL: DESCRIPTION AND PARAMETER ESTIMATION

Alternative approaches have been proposed to account for learning in games (see, e.g., Camerer, 2003 for an excellent review). They include evolutionary dynamics, various forms of reinforcement learning (McAllister, 1991; Roth & Erev, 1995; Sarin & Vahid, 2001), belief learning (Cheung & Friedman 1997; Fudenberg & Levine, 1998), learning direction theory (Selten & Stocker, 1986), Bayesian learning

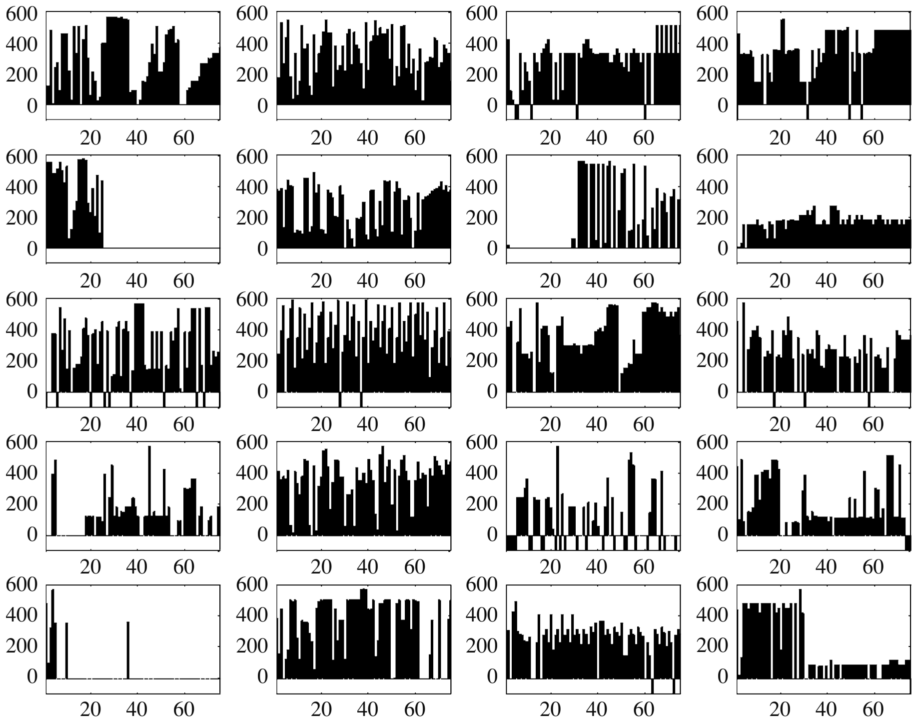

Figure 7. Individual decisions of all twenty subjects in Group 1 of Condition 1.

(Jordan, 1991), experience-weighted attraction (EWA) learning (Camerer & Ho, 1999), and rule learning (Stahl, 1996). Without making additional assumptions, these models are not directly applicable to our data.[1] We report below a simple learning model, which was constructed to account for the individual and aggregate patterns of our data reported above. This is clearly an *ad-hoc* model that does not have the generality of the approaches to learning mentioned above.

Basic Assumptions. The learning model uses a simple reinforcement learning mechanism to update arrival times based on historical play. It is derived from two primitive assumptions:

- Decisions to enter the queue are based on previous payoffs: as the agent's payoff on trial $t-1$ decreases, the agent is less likely to enter the queue.
- Once an agent has decided to enter the queue on trial t, its entry time is based on its entry times and payoffs on previous trials.

Both of these assumptions are consistent with the experimental data. Next, we describe a formal model that is derived from these assumptions.

Sketch of the Learning Model. The intuition underlying our learning algorithm is quite simple. On each trial t, the agent makes a decision either to enter the queue or not. If her payoff on trial $t - 1$ is high, then the agent enters with a higher probability than if the payoff was low. Put differently, the agents are more likely to *stay out* of the queue on a given trial if they did poorly on the previous trial. The agent's decision regarding *when* to enter the queue (conditional on her decision to enter) is based on her past decisions and the payoffs associated with those decisions. If an agent enters the queue at trial $t - 1$ and receives a good payoff, then she is likely to enter around that time again on trial t; on the other hand, if the agent receives a poor payoff for that entry time, then she is likely to change her entry time by quite a bit. Furthermore, if an increase (decrease) in arrival time consistently yields higher payoffs, then the agent is going to consistently increase (decrease) her arrival time. Increases (decreases) in arrival time that lead to poorer payoffs will cause the agent to decrease (increase) her arrival time. These learning mechanisms are formally specified in the following section.

Formal Specification of the Learning Model. Denote the entry time and payoff of agent i on trial t by A_t^i and π_t^i, respectively. If the queue is *entered*, then with probability $1 - \varepsilon$ entry times on the next trial are based on the following motion equations:

$$A_t^i = A_{t-1}^i + \begin{cases} \delta_t^i \eta_t^i \beta^i [(T_e - d) - A_{t-1}^i] & \text{if } \delta_t^i = +1 \\ \delta_t^i \eta_t^i \beta^i \,|A_{t-1}^i + (T_o - T_{\min})| & \text{if } \delta_t^i = -1 \end{cases}, \tag{1}$$

where

$$\delta_t^i = \begin{cases} +1 & \text{if } \pi_{t-1}^i \geq \pi_{t-2}^i \\ -1 & \text{if } \pi_{t-1}^i < \pi_{t-2}^i \end{cases}, \tag{2}$$

and

$$\eta_t^i = 1 - \exp[\tau^i(\pi_{t-1}^i - r)]. \tag{3}$$

With probability ε, A_t^i is sampled from a uniform probability distribution on the interval $[T_o - T_{\min}, T_e - d]$. (Without this "error" probability, the model produces individual subject results quite inconsistent with the individual subject experimental results.) The parameter β^i $(0 < \beta^i < 1)$ denotes the agent's *learning rate*, $T_{\min}$ is the earliest time the agent can enter the queue, τ^i is the agent's *payoff sensitivity*, and r is the payoff for completing service.

As for trial 1, by assumption A_1^i is sampled from a uniform discrete probability distribution defined on the interval $[T_o - T_{\min}, T_e - d]$, δ_t^i $(t = 1, 2)$ are sampled independently and with equal probability from the set $\{-1, +1\}$, and π_0^i is sampled with uniform probability from [0, r]. This initialization is conducted independently for each agent i. If the queue is *not entered* on trial t, then $A_t^i = A_{t-1}^i$. Thus, queue arrival time updates are always based on the most recently updated arrival time; arrival times are not updated during periods in which the agent does not enter the queue.

Decisions to enter the queue are made probabilistically; specifically, in the absence of group information (Conditions 1, 2P, and 3P), the probability of agent i entering the queue on trial t is given by

$$p_i^t = \exp[\lambda^i(\pi_{t-1}^i - r)]. \tag{4}$$

The parameter $\lambda^i > 0$ is the agent's *entry propensity*. Note that as λ^i approaches 0, the agent's entry probability goes to 1; and as λ^i goes to infinity, the entry probability goes to 0 (when, of course, $\pi_t^i < r$). The probability expressed in Eq. 4 is transformed in the Group Information Conditions (2G and 3G) as follows:

$$\tilde{p}_t^i = \begin{cases} p_t^i & \text{if } n_{t-1} = n_{cap} \\ \alpha_i p_t^i & \text{if } n_{t-1} > n_{cap} \\ p_t^i + \alpha_i(1 - p_t^i) & \text{if } n_{t-1} < n_{cap} \end{cases}. \tag{5}$$

where n_{cap} denotes the queue capacity. In Conditions 2 and 3, $n_{cap} = 20$ and $n_{cap} = 13$, respectively. The actual number of agents entering the queue on trial t is denoted by n_t. According to Eq. 5, entry probabilities are increased if the queue has too few entrants on the previous trial and are decreased if it has too many. The magnitude of the adjustment is determined by the parameter $0 < \alpha_i < 1$.

Model Parameter Estimation. To test the model's ability to capture the important properties of the experimental data, we first found *best fitting parameters* for the model using a grid search (brute force) algorithm. Goodness of fit was estimated by comparing the model's arrival time distributions to those from the experimental subjects.

Let C_T denote the proportion of arrival times less than or equal to T. Model fit was measured as the root-mean-square deviation of the model arrival time distribution from the subject's arrival time distribution:

$$RMSD = \left[\frac{1}{(T_{\max} - T_{\min} + 1)} \sum_{T=T_{\min}}^{T_{\max}} (C_T^M - C_T^D)^2\right]^{1/2}, \tag{6}$$

where C^M are the learning model cumulative arrival times and C^D are those of the experimental subjects. The proportion of non-entries is given by $1 - C_{T_e}$. Thus, optimal fitting involves finding the vector $V = (a, b, \tau, \lambda, \alpha)$ such that $RMSD$ is minimized, where a and b are the parameters of the beta distribution $B(a, b)$ from which the β^i are independently sampled for each simulated subject i. In the results reported here, for all simulated agents we assume that $\tau^i = \tau$ for all i ($i = 1, \ldots, N$), and likewise for λ^i and α^i. (A study of the model output suggested that allowing β^i to be a random variable, while making all other model parameters constant, was necessary to capture important properties of the experimental results. Allowing for all of the parameters to be random variables simply introduces too many parameters (as the distribution of random variables must be parameterized, which, in the case of, say, a beta distribution, introduces two distribution parameters for a single model parameter). It is our contention that the model results support this approach.) Since the agents only receive private information in Conditions 1, 2P, and 3P, α is constrained to equal 1 in these conditions. We fixed ε to be equal to 0.10 when we estimated all other parameters (the objective function was relatively flat with respect to ε, making estimating ε using monte carlo methods very difficult). Thus, we must estimate four parameters in Conditions 1, 2P, and 3P; all five parameters must be estimated in Conditions 2G and 3G. For each experimental condition, C^M was *estimated* for each V by aggregating the arrival times from 100 independent simulations of 75 trials of play of the queuing game with 20 agents. Since our objective function can only be estimated through simulation, one concern is that we might obtain inconsistent estimates of V; however, multiple replications of the grid search algorithm produced highly consistent results.

5. TESTING THE LEARNING MODEL

5.1. Condition 1

Aggregate Arrival Time Distributions. The cumulative arrival time distributions for the experimental subjects and the simulated learning agents, as well as the equilibrium cumulative arrival time distribution, are displayed in Fig. 4. With the exception of the aggregate arrival time at 8:00 (where the model under-predicts the probability of arrival), the model results closely agree with those of the human subjects.

Individual Arrival Times. Fig. 7 exhibits the individual arrival times of the 20 subjects in Group 1 of Condition 1. The decisions to stay out are represented by the downward ticks on the horizontal axis. Individual arrival time distributions for 20 *simulated* agents in Condition 1 are shown in Fig. 8. Observe that both the human subjects and simulated agents display heterogeneous arrival time behavior. Some subjects switch their arrival times quite often and quite dramatically, while others make less frequent and less dramatic switches. There is no simple way of telling which figure displays the individual arrival times of the genuine subjects and which of the simulated agents.

Switching Behavior. Fig. 9 shows the mean switching probabilities and mean switch magnitudes across trials for the human subjects on Condition 1. Here, a switch obtains on trial t when the subject (or simulated agent) enters on both trials $t - 1$

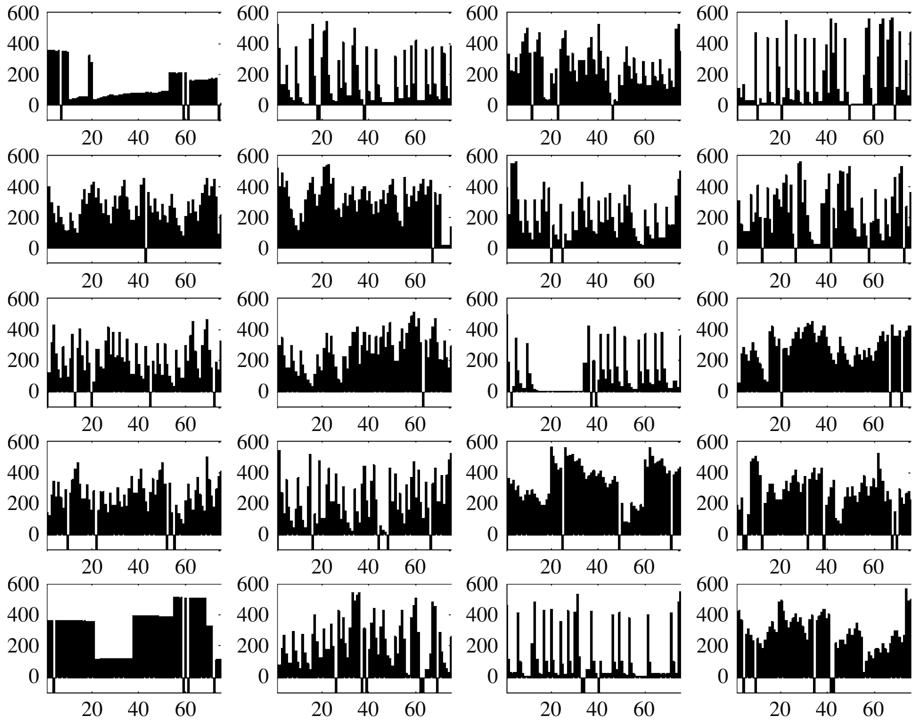

Figure 8. Individual decisions of twenty simulated agents in Condition 1.

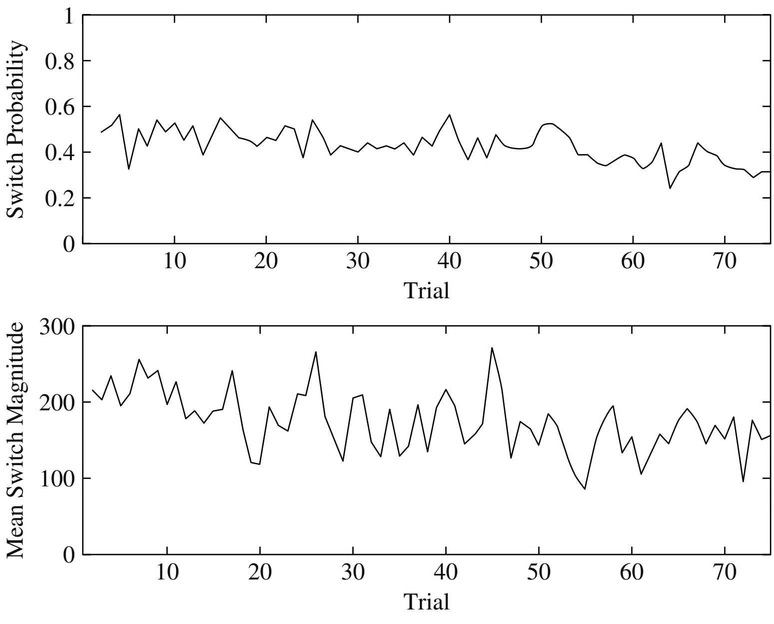

Figure 9. Switch probabilities and mean switch magnitudes across trials for all four experimental groups in Condition 1.

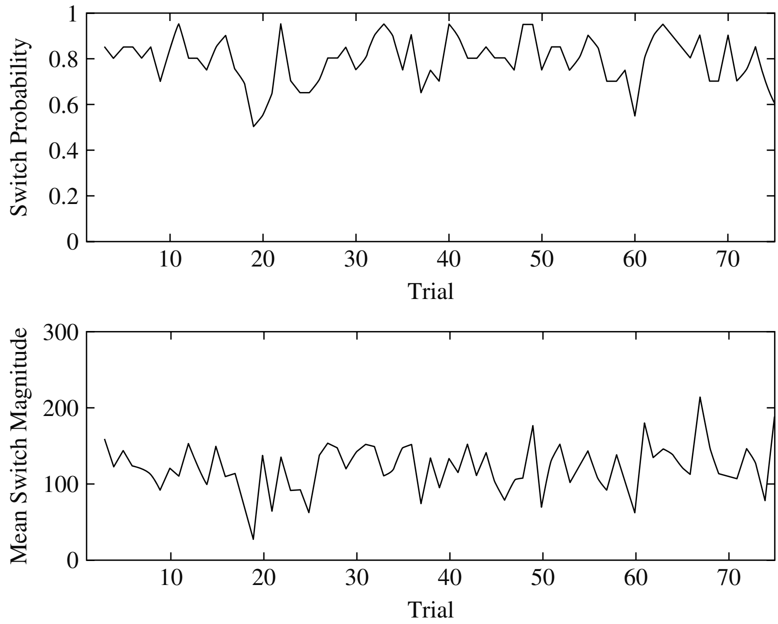

Figure 10. Switch probabilities and mean switch magnitudes across trials for four simulated groups in Condition 1.

and t and $A_{t-1} \neq A_t$. The magnitude of a switch is defined as the absolute difference between arrival times on trials $t - 1$ and t. The corresponding plot for the simulated agents, which is based on the best-fitting parameters shown in Fig. 4, is exhibited in Fig. 10. A comparison of Figs. 9 and 10 shows basically no change in the trend in mean switch probability and mean switch magnitude across trials for both the simulated and genuine subjects. However, the mean switch probabilities for the simulated agents consistently exceed the ones for the experimental subjects by more than 50%. Also, we observe that the mean switch magnitudes for the simulated subjects are slightly lower than those for the experimental subjects.

5.2. Conditions 2 and 3

Arrival Time Distributions. Figs. 5 and 6 display the cumulative arrival time distributions for Conditions 2 and 3, respectively. The distributions for the private outcome information (Conditions 2P and 3P) are displayed on the upper panels, and those for the group outcome information (Conditions 2G and 3G) on the bottom panels. The learning model results and the experimental data are in close agreement. In fact, the learning model accounts better for the results of Conditions 2 and 3 than Condition 1. The only notable discrepancy is in Condition 3P, where the model entry probability is about 0.05 greater than that of the human subjects. As the results for individual arrival time distributions, mean probability of switching, and mean magnitude of switching are similar to those in Condition 1, they are not exhibited here. Again, we observe a higher probability of switching and smaller mean switch magnitude in the simulated agents.

6. DISCUSSION AND CONCLUSION

RSPS and SPSR have studied experimentally how delay-averse subjects, who patronize the same service facility and choose their arrival times from a discrete set of time intervals simultaneously, seek service. Taking into account the actions of others, whose number is assumed to be commonly known, each self-interested subject attempts to maximize her net utility by arriving with as few other subjects as possible. She is also given the option of staying out of the queue on any particular trial. Using a repeated game design and several variants of the queueing game, RSPS and subsequently SPSR reported consistent patterns of behavior (arrival times and staying out decisions) that are accounted for successfully by the symmetric mixed-strategy equilibria for these variants, substantial individual differences in behavior, and learning trends across iterations of the stage game. Our major purpose has been to account for the major results of several different conditions by the same reinforcement-based learning model formulated at the individual level.

Our "bottom-to-top" approach to explain the dynamics of this repeated interaction calls for starting the analysis with a simple model that has as few parameters as possible, modify it, if necessary, in light of the discrepancies between theoretical and observed results, and then apply it to other sets of data. The focus of the present analysis has been on the distributions of arrival time on both the aggregate and individual levels. Although our learning model has been tailored for a class of queueing games with endogenous arrivals, it has some generality as it is designed to account for the results in five different conditions (1, 2P, 3P, 2G, 3G) that vary from one another on several dimensions.

The performance of the model is mixed. It accounts quite well for the aggregate distributions of arrival time in four of the five conditions. (The main exception is the aggregate arrival time at 8:00 in Condition 1.) For many learning models, this is the major criterion for assessing the model performance. The model also produces heterogeneous patterns of individual arrival times that are quite consistent with those of experimental subjects.

On the negative side, the learning model generates considerably more switches than observed in the data and somewhat smaller mean switch magnitude than observed in all the experimental conditions. Analysis of individual decisions in the studies by RSPS and SPSR shows that some subjects often enter the queue repeatedly at the same time, regardless of the outcomes on previous trials, possibly in an attempt to scare off other subjects or simply observe the pattern of entry without committing themselves to switch their arrival times. This kind of forward looking behavior, which is not captured by the learning model or any other reinforcement-based model in which a decision on trial t only depends on past decisions and outcomes, could be accounted for by increasing the complexity of the model. Although we only focus on testing a single learning model, our position is that in a final analysis the predictive power, utility, and generalizability of a learning model could better be assessed by comparing it to alternative models.

ACKNOWLEDGMENT

We gratefully acknowledge financial support by NSF Grant No. SES-0135811 to D. A. Seale and A. Rapoport and by a contract F49620-03-1-0377 from the AFOSR/ MURI to the Department of Industrial Engineering and the Department of Management and Policy at the University of Arizona.

NOTE

[1] We verified this for a Roth-Erev-type reinforcement-based learning model. With our implementation, we have been unable to reproduce most of the regularities we observe in the experimental data.

REFERENCES

Camerer, C. F. (2003). Behavioral game theory: Experiments on atrategic interaction, Princeton: Princeton University Press.

Camerer, C. F. and Ho, Teck (1999). "Experienced-weighted attraction learning in normal-form games." *Econometrica*, 67, 827–874.

Cheung, Y-W., and Friedman, D. (1997). "Individual learning in normal form games: Some laboratory results." *Games and Economic Behavior*, 25, 34–78.

Fudenberg, D. and Levine, D. (1998). *The Theory of Learning in Games*. Cambridge: Mass: MIT Press.

Hassin, R. and Haviv, M. (2003). *To Queue or Not to Queue: Equilibrium Behavior in Queueing Systems*. Boston: Kluwer Academic Press.

Jordan, J. S. (1991). "Bayesian learning in normal form games." *Games and Economic Behavior*, 3, 60–81.

Lariviere, M. A. and Mieghem, J. A. (2003). Strategically seeking service: How competition can guarantee Poisson arrivals. Northwestern University, Kellogg School of Business, unpublished manuscript.

McAllister, P. H. (1991). "Adaptive approaches to stochastic programming." *Annals of Operations Research*, 30, 45–62.

Rapoport, A., Stein, W. E., Parco, J. E., and Seale, D. A. (in press). "Strategic play in single-server queues with endogenously determined arrival times." *Journal of Economic Behavior and Organization*.

Roth, A. E. and Erev, I. (1995). "Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term." *Games and Economic Behavior*, 8, 164–212.

Sarin, R. and Vahid, F. (2001). "Predicting how people play games: A simple dynamic model of choice." *Games and Economic Behavior*, 34, 104–122.

Seale, D. A., Parco, J. E., Stein, W. E., and Rapoport, A. (2003). Joining a queue or staying out: Effects of information structure and service time on arrival and exit decisions. Department of Management and Policy, University of Arizona, unpublished manuscript.

Selten, R. and Stocker, R. (1986). "End behavior in sequences of finite Prisoner Dilemma's supergames: A learning theory approach." *Journal of Economic Behavior and Organization*, 7, 47–70.

Stahl, D. O. (1996). "Boundedly rational rule learning in a guessing game." *Games and Economic Behavior*, 16, 303–330.

Chapter 12

DECISION MAKING WITH NAÏVE ADVICE

Andrew Schotter
New York University

Abstract

In many of the decisions we make we rely on the advice of others who have preceded us. For example, before we buy a car, choose a dentist, choose a spouse, find a school for our children, sign on to a retirement plan, etc. we usually ask the advice of others who have experience with such decisions. The same is true when we make major financial decisions. Here people easily take advice from their fellow workers or relatives as to how to choose stock, balance a portfolio, or save for their child's education. Although some advice we get is from experts, most of the time we make our decisions relying only on the rather uninformed word-of-mouth advice we get from our friends or neighbors. We call this ?aive advice? In this paper I will outline a set of experimental results that indicate that word-of-mouth advice is a very powerful force in shaping the decisions that people make and tends to push those decisions in the direction of the predictions of the rational theory.

1. INTRODUCTION

In many of the decisions we make we rely on the advice of others who have preceded us. For example, before we buy a car, choose a dentist, choose a spouse, find a school for our children, sign on to a retirement plan, etc., we usually ask the advice of others who have experience with such decisions. The same is true when we make major financial decisions. Here people easily take advice from their fellow workers or relatives as to how to choose stock, balance a portfolio, or save for their child's education. Although some advice we get is from experts, most of the time we make our decisions relying only on the rather uninformed word-of-mouth advice we get from our friends or neighbors. We call this "naive advice".

Despite our everyday reliance on advice, economic theory has relatively little to say about it. Hence, there tends to be relatively little written in the decision theoretic or game theoretical literature about decision making with advice.

In this paper I outline a set of experimental results (see, Schotter and Sopher, 2003, 2004a, 2004b, Chaudhri, Schotter, and Sopher, 2002; Iyengar and Schotter, 2002; and Celen, Kariv, and Schotter, 2003) indicating that word-of-mouth advice is

A. Rapoport and R. Zwick (eds.), Experimental Business Research, Vol. II, 223–248.

a very powerful force in shaping the decisions that people make, and tends to push those decisions in the direction of the predictions of the rational theory. More precisely, I will demonstrate the following:

1) Laboratory subjects tend to follow the advice of naive advisors, i.e., advisors that are hardly any more expert in the task they are engaged in than they are.
2) This advice changes their behavior in the sense that subjects who play games or make decisions with naive advice play differently than those who play identical games without such advice.
3) The decisions made in games played with naive advice are closer to the predictions of economic theory than those made without it.
4) If given a choice between getting advice or the information upon which that advice was based, subjects tend to opt for the advice indicating a kind of under-confidence in their decision making abilities that is counter to the usual ego-centric bias or overconfidence observed by psychologists.
5) The reason why advice increases efficiency or rationality is that the process of giving or receiving advice forces decision makers to think about the problem they are facing in a way different from the way they would do so if no advice was offered.

In many of the experiments reported below, subjects engage in what are called "intergenerational games". In these games, a sequence of non-overlapping "generations" of players play a stage game for a finite number of periods and are then replaced by other agents who continue the game in their role for an identical length of time.[1] Players in generation t are allowed to observe the history of the game played by all (or some subset) of the generations who played it before them and can communicate with their successors in generation $t + 1$ and advise them on how they should behave. This advice is in two parts. First, in most of the experiments discussed below, subjects offer their successors a strategy to follow. After this they may write a free-form message giving the reasons why they are suggesting the strategy they are. These messages are a treasure trove of information about how these subjects are thinking the problem through. Because they have incentives to pass on truthful advice (they are paid 1/2 off what their successors earn), we feel confident that this advice is in earnest. Hence, when a generation t player is about to move she has both history and advice at her disposal. (Actually, we investigate three experimental treatments. In one that we call the *Baseline*, when generation t replaces generation $t - 1$, subjects are allowed to see the history of play of all previous generations and receive advice from their predecessors. This advice is almost always private between a generation $t - 1$ player and his progeny. In a second treatment called the *History-Only* treatment, subjects can see the entire history but receive no advice from their predecessors. Finally, in our third treatment called the *Advice-Only* treatment, subjects can receive advice but can only view the play of their immediate predecessor's generation). In addition, players care about the succeeding generation in the sense that each generation's payoff is a function not only of the payoffs achieved during their generation but also of the payoffs achieved by their children in the game that is played after they retire.[2] By comparing the play of subjects in these three treatments we can measure the impact of advice on behavior.

In the remainder of this paper we will survey the papers cited and use the result generated there to substantiate the statements made above.

2. THE IMPACT OF ADVICE

2.1. Ultimatum Games (Schotter and Sopher (2004a))

Consider an Ultimatum Game with a \$10 endowment played as an inter-generational game where each generation plays once and only once before it is retired. In our experiments we had 81, 79, and 66 generations play this game under the three treatments described above, respectively.

Since this game is played inter-generationally with each generation playing once and only once, when a Proposer arrives in the lab he sees on his computer screen an amount advised to be sent. A Receiver receives advice advising her what the minimum offer she should accept. Economic theory predicts that only a small amount, \$.01 say, will be sent.

2.2. Was Advice Followed?

2.2.1. Offer Behavior

Figures 1a and 1b display the amounts advised to be sent as well as the amounts actually sent by each generation in two treatments of our intergenerational Ultimatum game experiment – the Baseline treatment (where subjects can both receive advice and see the entire history of all generations before them) and the Advice-Only treatment (where subjects can receive advice but only see the history of their immediate predecessors).

As can easily be seen, by and large subjects simply sent the amount they were advised to send. Advice was followed in a very direct way.

2.2.2. Was Behavior Changed by Advice?

Advice had a significant impact on behavior. For example, while the mean amount offered in the Advice-Only experiment over the last 40 generations was 33.68, it was 43.90 in the History-Only treatment. Figures 2a–2b show the amounts offered by Proposer subjects in two experiments – the Advice-Only experiment (Treatment I), where advice was allowed but history eliminated, and the History-Only Experiment (Treatment II), where subjects could see the entire history but not see advice.

(a)

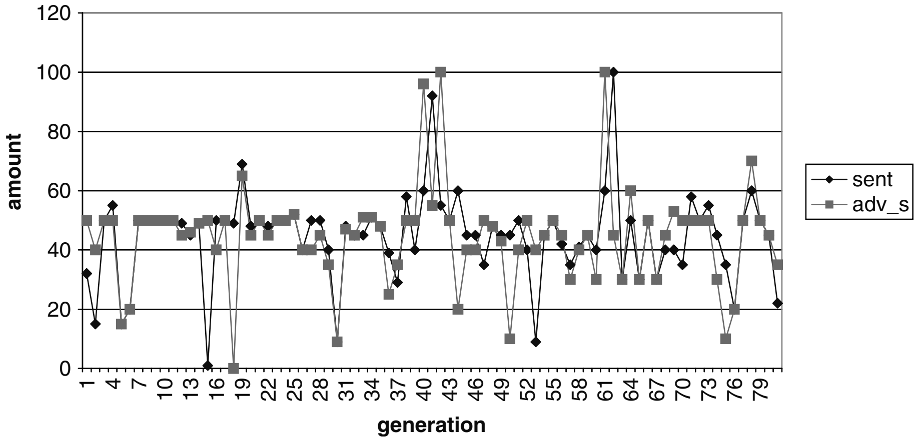

(b)

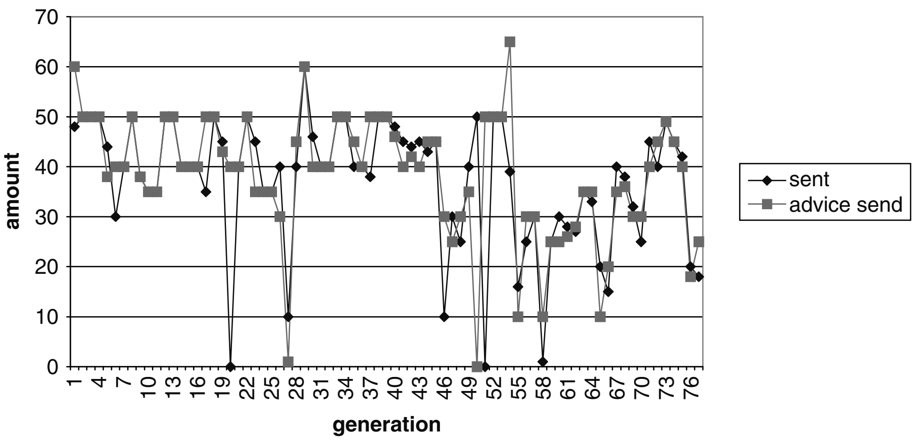

Figure 1. (a) Amount Sent and Advice: Baseline; (b) Amount Sent and Advice: Advice Only Treatment.

Note that the impact of advice is to truncate the right tail of the offer distribution and decrese the variance of offers made. In fact, while only 10% of the offers in the Advice-Only treatment were greater than 50, in the History-Only treatment 10% of the observations were above 80. A series of one-tailed F-tests supports this observation for binary comparisons between with the History-only treatment and the Baseline ($F_{\{(65,80)\}} = 2.16$, $p = .00$) and the History-only treatment and the Advice-only treatment ($F_{\{(65,76)\}} = 2.90$, $p = .00$). The same test found a difference between the variances of the Advice-only treatment and the Baseline at only the 10% level. These results indicate that history does not seem to supply a sufficient lesson for subjects to guide their behavior in a smooth and consistent manner.

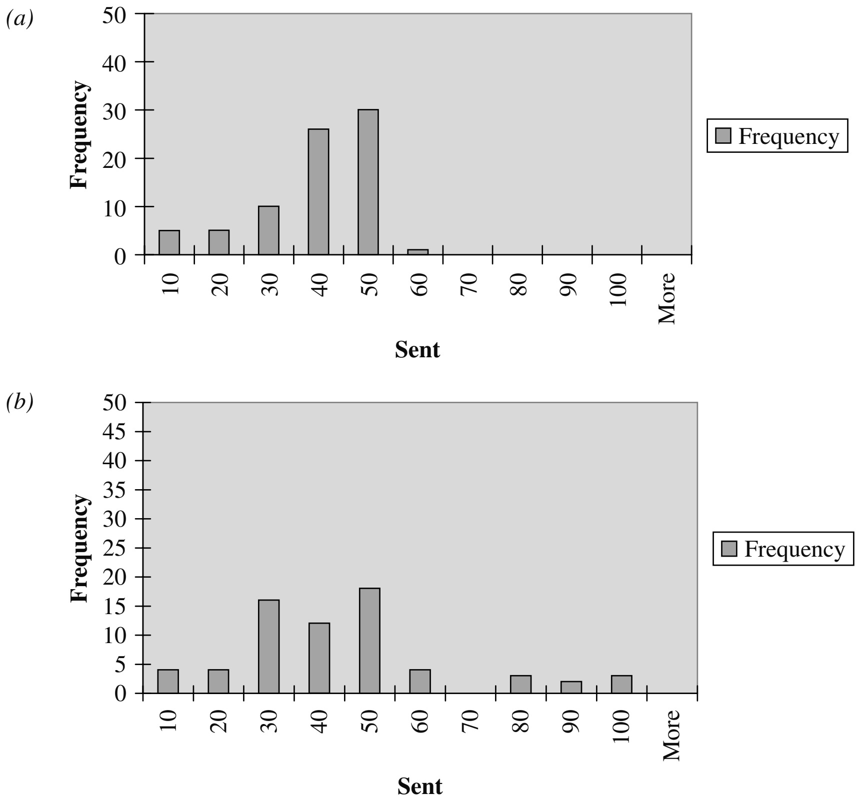

Figure 2. (a) Amount Sent Advice-Only Experiment Treatment I; (b) Amount Sent History-Only Treatment II.

2.2.3. Rejection Behavior

Rejection behavior is also affected by advice. Schotter and Sopher (2003a) used a logit model to estimate the probability of acceptance as a function of the amount sent of the following type:

$$\Pr(x \text{ accepted}) = e^{a+bx}/(1 + e^{a+bx}),$$

where x is the amount offered and the left hand variable is a $\{0, 1\}$; the variable taking a value of 1 if x is accepted and 0 otherwise.

The results of these estimations are presented in Figure 3 that plots the resulting estimated acceptance functions and superimposes them on the same graph.

Figure 3 shows that low offers are least likely to be accepted when only advice exists (the Advice-Only treatment) and most likely to be accepted when no advice is present but access to history is unlimited (the History-Only treatment). The Baseline, in which both treatments exist simultaneously, is in between. For example, while the

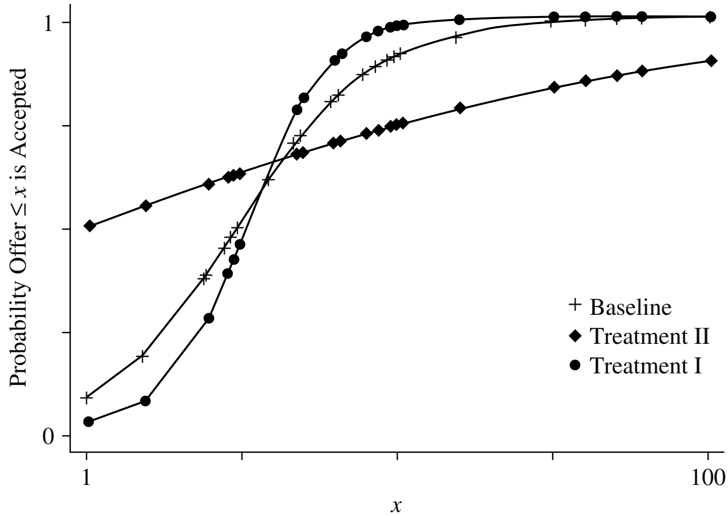

Figure 3. Acceptance Behavior.

probability that an offer of 10 is accepted is about 0.10 in the Advice-Only treatment, that probability increases to about 0.19 and 0.53 in the Baseline and History-Only treatments, respectively.

2.3. Coordination Conventions (Schotter and Sopher (2003)

Consider the following Battle of the Sexes game played in the lab as an intergenerational game:

Battle of the Sexes Game
Column Player

		1	2
Row Player	1	150, 50	0, 0
	2	0, 0	50, 150

2.3.1. Was Advice Followed?

In the Baseline treatment of our Battle of the Sexes game advice appears to be followed quite often but the degree to which it is followed varies depending on the state last period. On average, for the row players it is followed 68.75% of the time, while for the column player it was followed 70% of the time. When the last period state was (2, 2) (i.e., when in the last period the subjects coordinated on the (2, 2) equilibrium), row players followed the advice given to them 73.3% of the time while

Table 1. Advice Following when Advice and Best Response Differ and when They are the Same – Baseline Condition

Advice Differs from Best Response			
Row Follows	*Row Rejects*	*Column Follows*	*Column Rejects*
15	13	17	17
Advice Equals Best Response			
40	12	39	7

column subjects followed 86.6% of the time. When the last period state was the (1, 1) equilibrium, column subjects chose to follow it only 37.5% of the time while row player adhered 68% of the time.

In these experiments, we measured the beliefs of each generation concerning their expectations of what strategies they expect their opponent to choose. We did this using a proper scoring rule, and this enabled us to define what a subject's best response was to those beliefs. Since in some instances the advice offered to subjects was counter to their best response action, we can measure the relative strength of advice by comparing how often the subjects chose one over the other.

When advice and best responses differ, subjects are about as likely to follow the dictates of their best responses as they are those of the advice they are given. Consider Table 1 that presents data from our Baseline experiment. .

As we can see, for the row players there were 28 instances where the best response prescription was different than the advice given, and of those 28 instances the advice was followed 15 times. For the column players there were 34 such instances, and in 17 of them the column player chose to follow advice and not to best respond. When advice supported the best response of the subject, we see that it was adhered to more frequently (79 out of 98 times).

These results are striking since the beliefs we measured were the player's posterior beliefs after they had both seen the advice given to them and the history of play before them. Hence, our beliefs should have included any informational content contained in the advice subjects were given, yet half of the time they still persisted in making a choice that was inconsistent with their best response. Since advice in this experiment was a type of private cheap talk based on little more information than the next generation already possesses (the only informational difference between a generation t and generation $t + 1$ player is the fact that the generation t player happened to have played the game once and received advice from his or her predecessor which our generation $t + 1$ player did not see directly), it is surprising that it was listened to at all.

2.3.2. Was Behavior Changed by Advice?

One puzzle that arises from our Battle of the Sexes experiments is the following. While in the Baseline we observe equilibrium outcomes 58% of the time (47 out of 81 generations), when we eliminate advice, as we do in History-Only Treatment, we observe coordination in only 29% of the time (19 out of 66 generations). When we allow advice but remove history, the Advice-Only treatment, coordination is restored and occurs 49% of the time (39 out of 81 generations).

These results raise what we call the "Advice Puzzle" which is composed of two parts. Part 1 is the question of why subjects would follow the advice of someone whose information set contains virtually the same information as theirs. In fact, as stated above, the only difference between the information sets of parents and children in our Baseline condition is the advice that predecessors received from their predecessors.

Part 2 of the Advice Puzzle is that despite the fact that advice is private and not common knowledge cheap talk, as in Cooper, Dejong, Forsythe and Ross (1989), it appears to aid coordination in the sense that the amount of equilibrium occurrences in our Baseline (58%) and Advice-Only treatment (49%) where advice was present is far greater than that of History-Only treatment (29%) where no advice was present. While it is known that one-way communication in the form of cheap talk can increase coordination in Battle of the Sexes games (see Cooper et al. (1989)), and that two-way cheap talk can help in other games, (see Cooper, Dejong, Forsythe and Ross (1992)), how private communication of the type seen in our experiment works is an unsolved puzzle for us.

Finally, note that the desire of subjects to follow advice has some of the characteristics of an information cascade since in many cases subjects are not relying on their own beliefs, which are based on the information contained in the history of the game, but are instead following the advice given to them by their predecessor who is as just about much a neophyte as they are.

2.4. Trust Games (Schotter and Sopher (2004b))

The particular trust game that we consider, first investigated by Berg, McCabe and Dickhaut (1995), is the following. Player 1 moves first and can send Player 2 any amount of money x in [0, 100] or keep 100 for herself. Once x is determined, it is multiplied by 3 and the amount $3x$ is received by Player 2. Player 2 can then decide how much of the $3x$ received, y, to send back to Player 1. The payoffs for the players are then $100 - x + y$ for Player 1 and $3x - y$ for Player 2. Note that this game is a game of trust since Player 1, by sending nothing, can elect to get a safe payoff for himself of 100. But if he sends any amount x to Player 2, he places his fate in Player 2's hands and must trust him to reciprocate and send back at least x to compensate him for his act of trust. Hence, Player 2 is trustworthy if he sends back an amount $y \geq x$ and is not trustworthy, otherwise.

We played this game of trust in an inter-generational setting where a game is played by a pair of players who subsequently are replaced by another pair, each replacement being a "descendent" of one of the original players and able to receive advice from their predecessor on how to play of the game. We analyze the impact of this inter-generational advice on behavior.

What we find is consistent with the pattern of results reported above for the Ultimatum and Battle of the Sexes games with some, perhaps significant, differences.

2.5. Do Subjects Follow Advice and Does the Presence of Advice Change Behavior?

2.5.1. Sender Behavior

As we can see by observing Figures 4a and 4b, there appears to be a close qualitative relationship between advice given by Senders and the amounts sent by their successors. To the extent that subjects did not follow the advice of their predecessors, they did so by sending more than suggested and not less. Looking at Figures 4a and 4b we see that while the gyrations of the time series of amounts sent tends to track that of the advice time series, it also tends to be greater than it most of the time – subjects send more than advised. (We will see a similar result in the last section of this paper as well).

Despite the fact that subjects tend to reject the advice of their predecessors and send more than suggested, it is still true that when compared to the History-Only Experiment less is sent when advice is present. In other words, advice is trust decreasing. This can easily be seen in Figures 5a–5c, which present histograms of the amount sent in each treatment.

Note that in all treatments the amount sent is substantially above the zero predictions of the static sub-game perfect Nash equilibrium prediction. For example, in all of our treatment over 82% of the subjects send something positive. In the Baseline, 50% send 15 or more while in the Advice-only and History-only Treatments 50% send 20 or more and 40 or more, respectively. The presence of advice has a dramatic impact on sending behavior, however. As we can see in Figures 5a–5c, the amount sent is substantially higher in the History-Only Treatment than in either the Baseline or Advice-Only Treatments. For example, the mean (median) amount sent in the Baseline and Advice-only Treatments, respectively, is 25.94 (15) and 28.10 (25), while in the History-Only Treatment, where there is no advice, it was 40.18 (30). A set of two-sample Wilcoxon rank-sum tests indicate that while there is no significant difference between the samples of Baseline and Advice-Only Treatment offers (z statistic -1.24, p-value .22) , a significant difference did exist between the amounts sent in the History-Only Treatment and both the Baseline (z statistic -3.03, p-value .00) and Advice-only Treatments (z statistic -2.13, p-value .03).

In addition, while the inter-quartile range of offers in the Baseline and Advice-Only Treatments were 1–40 and 5–40 respectively, the same range was 15–55 in the History-Only Treatment. Another measure of trust can be gleaned from the upper end of the offer distribution. For instance, 10% of all offers in the Baseline and the

(a)

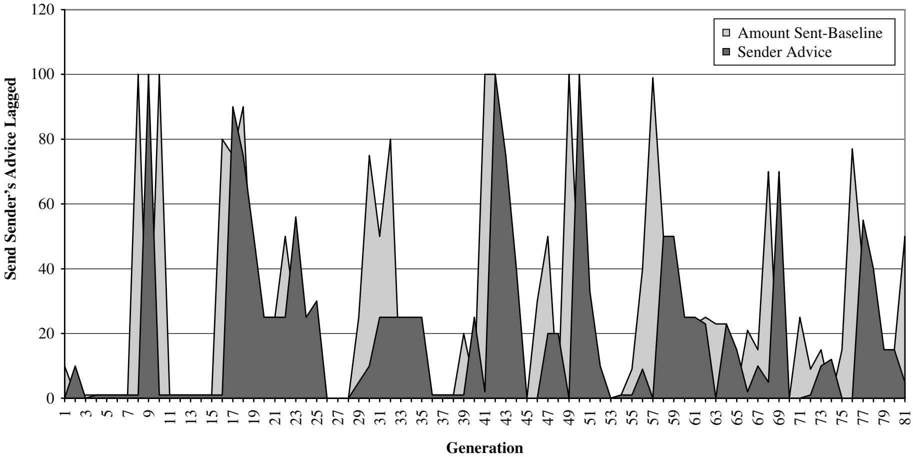

(b)

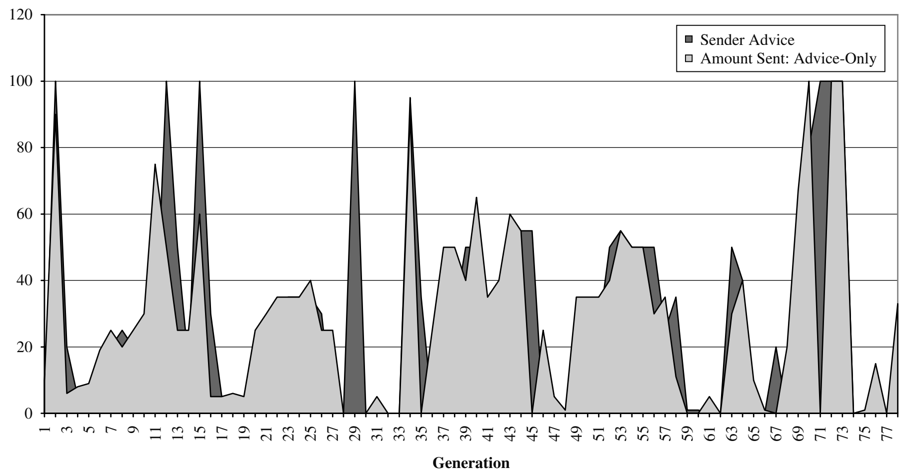

Figure 4. (a) Amount Sent and Amount Advised to Be Sent: Baseline Treatment; (b) Amount Sent and Amount Advised to be Sent: Advice-Only Treatment.

Advice-Only Treatment experiments were greater than 80 and 65 respectively, while 10% of all offers in the History-Only Treatment were equal to 100 indicating an extreme willingness to "risk it all". Finally, to demonstrate the impact of advice on amounts sent (as) we ran a regression of "as" on a {0, 1} dummy variable (D) depicting whether or not advice was allowed in the experiment generating the observation. According to this regression, we again observe a significant and negative relationship between the presence of advice and the amount sent. On the basis of

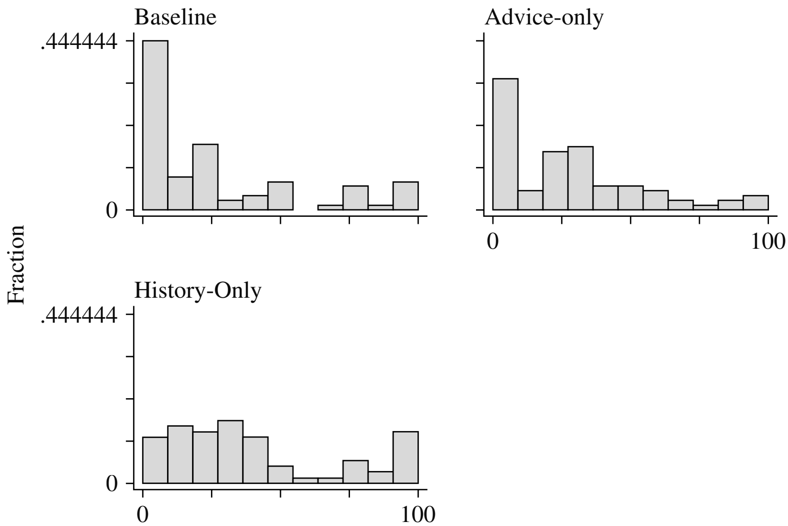

Figures 5a–5c. Amount Sent in Trust Game by Treatment.

these results we conclude that advice lowers the amount of trust in this game by lowering the amount of money sent.[3]

2.6. Was Receiver Behavior Changed By Advice? The Impact of Advice on Trustworthiness

Trustworthiness in these experiments is measured by how much of the amount sent is returned by the Receiver. When we look at the data from this experiment we see that while advice made the Senders less trusting, it made the Receivers more trustworthy. In short, while Receiver subjects tend to return 8.63 units less than they receive in the Baseline experiment, and 9.24 units less in the Advice-Only experiment, in the History-Only experiments they return, on average, 16.15 units less. The explanation, we believe, involves a small bit of anchoring and adjusting. In both the Sender and Receiver cases, subjects take the advice they are given and adjust from the suggested amounts. In the case of Senders, we know that subjects with no advice, the History-Only subjects, tend to send more than subjects in the Baseline and Advice-Only experiments are advised to. Hence, even though they send more than suggested, they ultimately send less than their non-advised History-Only counterparts. They use the advice as an anchor and adjust upwards. For the Receivers the effect is the opposite. While in the History-Only experiment sending back zero might be the natural anchor from which subjects adjust upward, in the Advice-Only and Baseline experiments the non-zero advice offered by one's advisor seems to function as the anchor from which subjects adjust upward. The net result is a higher amount of observed trustworthiness.

3. WOULD PEOPLE RATHER HAVE ADVICE OR DATA? (CELEN, KARIV, AND SCHOTTER, 2003)

In recent years, a great deal of attention has been paid to the problem of social learning. In the literature associated with this problem it is assumed that people learn by observing either all of or a subset of the actions of those who have gone before them.[4] They use these actions to update their beliefs about the payoff-relevant state of the world and then take an action that is optimal given those beliefs. Using this approach a great deal has been learned about how and why people follow their predecessors, or herd, and how informational cascades develop.

The odd aspect of the social learning literature as just described is that it is not very social. In the real world, while people learn by observing the actions of others, they also learn from their advice. For example, as stated in the introduction, people choose restaurants not only by viewing which of them are popular, but also by being advised to do so. People choose doctors not by viewing how crowded their waiting rooms are, but by asking advice about whom to go to, and so on. Thus, social learning tends to be far more social than economists depict it.

In the standard social learning situation, decision makers make their choices in sequence with one decision maker following the other. Typically, they are allowed to see what their predecessors have chosen after each of them receives an independent signal about the payoff-relevant state of the world. In CKS 2003, however, we allow decision makers, before they make a choice, to choose whether to observe the actions of those who went before or get advice from them as to what they should do. Which information do you think would be preferred? One would think that what you decide will depend on your estimate of your abilities as a decision maker compared to those of the advice giver as well as the informativeness of the data you might expect to process yourself.

To get at this question, Celen, Kariv and Schotter (2003) (CKS) investigated a social learning experiment with a design that differed slightly from the intergenerational game experiments described above. In this experiment eight subjects were brought into a lab and took decisions sequentially in a random order. A round started by having the computer randomly select eight numbers from the set of real numbers [−10, 10]. The numbers selected in each round were independent of each other and of the numbers selected in any of the other rounds. Each subject was informed only of the number corresponding to her turn to move. The value of this number was her private signal. In practice, subjects observed their signals up to two decimal points.

The task of subjects in the experiment was to choose one of two decisions labeled A and B. Decision A was the correct decision to make if the sum of the eight private signals was positive, while B was correct if the sum of the private signals was negative. A correct decision earned $2 while an incorrect one earned $0. This problem was repeated 15 times with each group of 8 decision makers each receiving a new and random place in the line of decision makers in each round.

Table 2. Agreement and Contrariness in Action-Only and Advice-Only Experiments

	Concurring	*Neutral*	*Contrary*
Action	44.2%	16.6%	39.2%
Advice	74.1%	9.1%	16.8%

CKS used three treatments that differed in the information they allowed subjects to have. In one treatment (the Action-Only treatment), subjects could see the decision made by their predecessor in the line of decision makers (so the fifth decision maker could see the decision of the fourth etc.) but no others, and could not receive any advice from their predecessors. In another treatment (the Advice-Only treatment), subjects (except for the first one) could receive advice from their predecessor telling them to either choose A of B. In the final treatment (the Advice-Plus-Action treatment), subjects could see both the decision their predecessor made and receive advice form him or her. Subject payoffs were equal to the sum of their payoffs over the 15 rounds in the experiment plus the sum of what their successors earned, so that each subject had an incentive to leave good advice. This design clearly makes the social learning problem more "social" by including elements of advice and word-of-mouth learning.

The final feature of the experimental design, and the one that distinguishes it from other social learning experiments, was that subjects did not directly choose a decision A or B but rather set a cut off level between −10 and 10 (a cutoff). Once this cutoff was typed into the computer it took action A for the decision maker if her signal was above the cutoff specified and action B if it was not.

This design can help us answer the question stated above; would people prefer to have advice or information. For example, Table 2 compares the actions of subjects who can only see the actions chosen by their immediate predecessor to those who cannot see what they have done, but can receive an advised action. I have broken down the actions of subjects into those actions which agree (concurring decisions), with the action or advice of the predecessor, those where the actions disagree (contrary decisions) and those where the actions neither agree or disagree with the actions or advice of one's predecessor (such actions are possible in this experiment since the subject can always set a zero cutoff which allows his to choose A or B with equal probability). By "agree" we mean that the subject sets a negative cutoff when he is told or observes the A action and sets a positive cutoff when he is told or observes the B action.

This table shows that subjects take actions that agree with the advice they receive 74.1% of the time yet copy the actions of their predecessors only 44.2% of the time. Actions disagree with advice only 16.8% of the time as compared with 39.2% for the experiment where actions only could be seen.

Table 3. Advice Taking in the Action-Plus Advice Experiment

	Successor	*Choose A*	*Choose B*	
Predecessor		Cutoff (−)	Cutoff (+)	Cutoff = 0
Action A/Advice B		13 (15.66%)	33 (39.76%)	6 (7.23%)
Action B/Advice A		17 (20.48%)	7 (8.43%)	7 (8.43%)

The Action-Plus-Advice experiment provides us with an extremely good opportunity to try separating the impact of advice and action on behavior. The reason is that in a number of situations subjects were faced with advice that was different from the action taken by the subject in the previous round. For example, in the Action-Plus-Advice experiment 83 out of the 525 decisions excluding the first decision turn (15.8 percent) were made under circumstances where the advice offered was different from the action observed in the previous period. If when these situations occurred, subjects chose to follow the advice of their predecessors rather than copying their action, we would interpret this as indicating that advice was more influential than action.

To pursue this line of inquiry, consider the choice of a negative cutoff as indicating a preference for the A choice and the choice of a positive cutoff as a preference for the B choice. If the advice and action of a predecessor subject differ, then two cases can be observed. The predecessor chooses A and advises B or the predecessor chooses B and advises A. Based on either of these occurring, the successor subject could choose to set either a negative cutoff (a higher probability of taking action A) or a positive one (a higher probability of taking action B). This defines four contingencies as depicted in Table 3.

Table 3 shows that when the advice and action of one's predecessor differ, successors are far more likely to choose an action consistent with the received advice than the observed action. For example, in 60.2 percent of the cases where the advice offered differs from the action, subjects chose to follow the advice they received rather than imitate their predecessor's action, while only 24.1 percent of the time they imitated the action taken, and 15.7 percent of the time they were neutral and choose a cutoff zero.

Table 3 looks at behavior when the advice offered by a subject's predecessor differs from the action she took. But we might also ask whether getting advice that is consistent with the action taken by one's predecessor makes a subject more likely to follow it and if so more likely to set a more extreme cutoff indicating stronger agreement. A priori we would expect this to be the case since when advice agrees with a predecessors' action we should expect a subject to view it as more compelling. Consider Table 4.

Table 4 supports our conjecture. Subjects are, indeed, more likely to follow advice (as indicated by the sign of their cutoff) when it is backed up by action. Note

Table 4. Decision Conformity with Advice and Action

	Action Taken		
	Concurring	*Neutral*	*Contrary*
Action-Only	44.2%	16.6%	39.2%
Advice-Only	74.1%	9.1%	16.8%
Action–+–Advice	84.2%	7.0%	8.8%

that if a subject is told to follow an action by a predecessor who took that action himself, such a recommendation is followed 84.2 percent of the time, while such advice is followed only 74.1 percent of the time in the Advice-Only experiment. When just the action is observed, it is imitated only 44.2 percent of the time. So it should be clear that a predecessor who does as she says is seen as being more believable than one whose advice cannot be backed up by action. Ironically, when a subject follows a piece of advice that is backed up by the actions of one's predecessor, the cutoff he sets is not significantly different than the one set by a subject in the Advice-Only experiment who also followed advice. Hence, it appears that while seeing actions support advice increases the probability of following the advice offered, the strength of conviction in the advice is not different from that in the Advice-Only Experiment.

3.1. Does Advice Increase Efficiency?

Probably the most important question that we can ask about the impact of advice on social learning is whether the presence of advice increases the welfare of subjects over and above what it would be without it. In answering this question, we will have to examine the impact that advice has on herding and cascade behavior of subjects since one way that advice affects behavior is through its propensity to cause subjects to herd with greater frequency than they would in its absence.

To begin, consider Table 5, which presents a summary our four experiments. It is clear that the mean payoffs of our subjects were highest in those experiments where advice was present.

As we see, while earnings for taking the correct action in the Action-Only experiment averaged \$18.8 they average \$23.3 and \$21.8 for the Action-Plus-Advice and Advice-Only experiments. These increases represent increases of 24.3 percent and 16.4 percent, respectively. In the Perfect-Information experiments of Celen and Kariv (2003) where subjects could see the entire history of actions before setting their cutoff values (but did not receive advice), earnings averaged \$22.0 indicating that advice with imperfect information is approximately as efficient as perfect information without advice. A set of binary Wilcoxon tests indicates that there is a

Table 5. Efficiency and Herding in Social Learning Experiments

	Action-Only	*Advice-Only*	*Action-Advice Information*	*Perfect*
Earnings	$18.8	$21.8	$23	$22
Herds*	8	25	36	27
% of Herds+	10.7	33.3	48	36
Cascades	18	24	21	26
% of cascades+	24	32	28	34.7

* Herds of at least five subjects
\+ Out of all 525 decision points excluding the first decision turn.

significant difference between the sample of subject payoffs in the Action-Only experiment and all other experiments at the 5 percent level of significance. It also indicates that no difference exists between the payoffs of subjects in the Perfect-Information experiment and any of those with advice, substantiating our conclusions that the presence of advice seems to be a substitute for the extra information contained in the perfect information experiment.

4. HERD BEHAVIOR AND INFORMATIONAL CASCADES

One of the main reasons why advice increases the payoffs and hence the welfare of our subjects is that it has a dramatic impact on our subjects' inclination to herd. We identify a subject who engages in cascade behavior as one who reports a cutoff of −$10 or $10, and thus takes either action A or B no matter what private signal she receives. In contrast, a subject who joins a herd but does engage in cascade behavior is one whose cutoff is in the open interval (−10, 10), indicating that there are some signals that can lead her to choose action A, some to choose B, but when her private signal is realized she will act as her predecessors did. Finally, we say that a *cascade* occurs when beginning with some subject all others thereafter follow cascade behavior, and *herd behavior* occurs when, beginning with some subject, all take the same action.

4.1. Herd behavior

While in our Action-Only experiments we observed herding of at least five subjects in only 8 of the 75 rounds (10.7 percent), in the Advice-Only and Action-Plus-Advice sessions herding occurred in 25 (33.3 percent) and 36 (48.0 percent) rounds respectively. Moreover, in the Action-Plus-Advice experiment herd behavior developed even more frequently than in the Perfect Information experiments of Celen and Kariv (2003) where subjects can see all the decisions made by all of their predecessors before making their choice, we found that herding was the outcome in 27 of the 75 rounds (36.0 percent). Finally, the frequency in which herd behavior occurs in the Action-Plus-Advice experiment compares favorably to the 47 percent predicted by the theory.

Obviously, two conditions must be met if advice is going to be welfare increasing. First, the advice must be correct and second it must be followed. Miraculously, in these experiments, both conditions seemed to have been met. In the Advice-Only experiments, whenever herd behavior arises all of the advice given was consistent with the action herded upon. In the Action-Plus-Advice experiments, this was not the case in only 5 of the 36 herds. In other words, when herds occurred those who herded tended to follow the advice given. More remarkably, in all experiments all herds turned out to be on the correct decision. This result is of a particular interest since one of the original concerns of the social learning literature was that herds and cascades might support or reinforce inefficient choices. Following Anderson and Holt (1997), these fears were supported by the results of many laboratory experiments.

To sum up, our results on herd behavior indicate that advice is a strong force in the creation of uniform social behavior and is welfare increasing.

4.2. Informational cascades

While all cascades must be herds, the opposite is certainly not true. Our experiment is uniquely designed to distinguish between the occurrence of cascades and herds since we are able to observe subjects' cutoffs that are typically unobservable. Surprisingly, advice did not have a significant impact on the rate of occurrence of information cascades. In the Action-Only experiments, cascades, in the sense that from some subject on all acted irrespective of the content of their private signals by setting either -10 or 10 as their cutoffs, were observed in 18 rounds (24.0 percent), whereas in the Advice-Only and Action-Plus-Advice experiments cascades formed in 24 (32.0 percent) and 21 (28.0 percent) rounds, respectively.

In summary, it appears that in this informational setting words speak louder than actions in the sense that subjects are more likely to follow the advice of their predecessors to take specific actions than they are to copy their behavior.

5. WHY FOLLOW ADVICE? (IYENGAR AND SCHOTTER (2002))

Our results above lead us to question why advice should be so beneficial. Why should people give better advice to their successors than they gave to themselves? An experiment conducted by Iyengar and Schotter (2002) attempts to answer this question.

In this experiment, subjects had to choose a number, e, between 0 and 100 called their decision number. They were told that they were playing against a computerized partner who would always choose the number 37. After this number is chosen a random number is independently generated from a uniform distribution over the interval $[-a, +a]$ for both the subject and his computerized opponent. These numbers (the decision number and the random number) are then added together and a "total number" is defined for each of the real and computerized players. Payoffs are determined by comparing the total numbers of the real and computerized subjects, and awarding the real player a fixed payoff of M if her total is larger than that of the computerized opponent. If her total number is smaller, then she receives a payoff of m, $m < M$. The cost of the decision number chosen is given by a convex function $c(e) = e^2/r$, where r is a constant. This amount is then subtracted from these fixed payments to determine a subject's final payoff. Hence, there is a trade-off in these experiments in the choice of decision numbers: higher numbers generate a higher probability of winning the big prize but at the same time also imply a higher decision cost. By letting $r = 500$, $a = 40$, $M = 29$ and $m = 17.2$, and holding the computerized player's choice fixed at 37, our subjects face a rather simple decision problem with a quadratic payoff function whose peak is at 37.

This task was used by Merlo and Schotter (1999, 2003) to test the impact of information on learning. They used what they called a "surprise quiz" method to test how well subjects learned the task put in front of them. In these experiments, subjects performed the exact task as described above 75 times and received payoffs each period. When the 75 rounds were over they were surprised and told that they would play the game once more but this time the stakes were multiplied by 75 so that they could earn for this one trial an amount equal to the sum of what they earned in all of the previous 75 rounds. Their choice in this high-stakes round should be a sufficient statistic for all that they have earned in the previous 75 rounds since the only way they can maximize their earnings in this round is by choosing that decision number which they feel is best. It is by comparing behavior in surprise quizzes that we can investigate the impact of various treatments on learning.

Merlo and Schotter (1999, 2003) performed this experiment under a number of different conditions. In one they simply had one subject perform the experiment exactly as described above and make a surprise quiz choice after the 75 rounds were done. In another experiment, the subject (which we call the doer) performed the experiment with another subject (the observer) silently watching what he did over his shoulder. After the 75 rounds were over we took the observer out of the room, he was told that he would do another experiment that was related to what the doer was doing but the specifics were not mentioned, and performed a surprise quiz on the doer. We then paid the doers and let them go and returned the observers to the lab where we announced that they were now to do a one-shot surprise quiz for 75 times the stakes of the doers they just watched.

Iyengar and Schotter (2002) repeated this experiment but in a slightly different manner. Instead of having one subject do the experiment alone, they sat another "advisor" subjects next to him or her at the computer. This advisor (type-P subject) makes written suggestions to the subject doing the experiment as to what he or she thinks is the best choice for that round. The chooser (type-A subject) is free to

Table 6. Experimental Design

	Treatment
Merlo and Schotter 1999	Decision maker alone
Merlo and Schotter 2003	Decision maker + overlooking subject
Iyengar and Schotter (2002)	Decision maker + Advisor

Table 7. Mean Surprise-Quiz Choices

	Mean (Median) Surpise-Quiz Choice		
Merlo-Schotter–1999		51.33	(50)
Merlo-Schotter–2003	Doer:	51.06	(50)
	Observer:	40.65	(37)
Iyengar-Schotter–2002	Doer: (No Cost)	31.20	(39.5)
	Advisor:	43.35	(43)
	Doer (Costly Advice):	36.1	(37)
	Advisor:	33.52	(38)

follow this advice or not but in one treatment is penalized for not doing so with a quadratic penalty function based on the difference between the action chosen and the action advised. In another treatment no penalty is assessed for not following advice (it is simply cheap talk). The advisor's payoff is equal to 3/4's of that of his advisee. After the initial 75 round experiment run in this manner, both he adviser and advisee are separated in given surprise quizzes. These experiments are summarized in Table 6 below:

It appears as if the process of giving advice and receiving greatly enhances the decision-making abilities of the Iyenga and Schotter subjects. Table 7 presents the mean choices of our subjects in the surprise quiz rounds in each of our four treatments.

In the MS (1999) experiment subjects performed our task alone without either an advisor or a spectator looking over their shoulder. In that experiment, as was true in the Merlo-Schotter (2003) experiment where subjects did the experiment with a spectator looking over their shoulder, subjects appeared to fail to learn where the peak of the payoff function was. In those experiments the mean and median surprise quiz choices of subjects who did the experiment for 75 rounds was 51 and 50, respectively.

The interesting result found in the MS (2003) paper was that while those who did the experiment failed to learn, those who watched them did quite well having a mean and median surprise-quiz choice of 40 and 37 respectively. In other words, watching was a better learning experience than doing; hence the title "Leaning by Not Doing". A median test rejects the hypothesis that 37 is the median of either the MS (2003) or MS (1999) doer choices but fails to reject the hypothesis that 37 was the median choice of observers in MS (2003). The message of these papers is that people fail to learn appropriately when they repeat experiments in which they receive payoffs after each period and the task is repeated often.

Iyengar and Schotter (2002), however, reported a remarkable result, namely, that the process of giving advice enhances the learning ability of both type-A and type-P subjects. For example, in the Costly-Advice experiment the mean and median choices of the type-A and type-P subjects were 36.1 and 37 for the type-A subjects, and 33.5 and 38 for the type-P subjects, respectively. Neither of these medians is significantly different from 37 using a median test at any meaningful level of significance (($p \leq 1$) for the type-A agents and ($p \leq .648$) for the type-P subjects). For the No-Cost Advice experiment, the situation is slightly different. Here the mean and median choices of the type-A and type-P subjects were 31.28 and 39.5 for the type-A and 43.35 and 43 for the type-P subjects. Only the type-A subjects had a median that was not significantly different from 37. For the type-P subjects we had to reject that hypothesis at the 2% level.

However, while this might indicate that advisors whose advice was ignored did not learn as well as those who received this advice (which they were at liberty to ignore), the type-P advisors still learned better than those in the MS (1999) and (2003) experiment who actually did the experiments. For example, a Wilcoxon test indicates that we can reject the hypothesis that the sample of type-P surprise quiz rounds came from the same population as either those of the subjects in MS (1999) ($p \approx 0$) or the doers in MS (2003) experiment ($p \approx 0$).

Another feature of the learning experience subjects have when advice is given is that advice seems to diminish the number of subjects who choose dominated strategies. For example, in the surprise quiz rounds of all four experiments (where there are no disagreement costs) any choice of 65 or more is dominated by choosing 0. While in MS (1999) 10 out of 24 subjects chose a dominated strategy in their surprise-quiz round and in MS (2003) 9 out of 31 doers did so, in our No-Cost Advice experiment only 1 type-P subject and 1 type-A subject made dominated choices. In other words, there were only 2 out of 28 such choices. For the Costly Advice experiment, the results are the same. Only 1 type-P subject and no type-A subjects made surprise-quiz choices strictly greater than 65. This is a very strong difference indicating that these subjects clearly learned some minimum lesson that seemed not to be learned by others in the No-Advice treatments.

The punch line then is that learning is fostered when advice is given even if there is no cost for ignoring it. Those who learn well are both the people who give the advice and those who receive it.

It is relevant to point out here that it appears that subjects have a hard time learning in environments where decisions are repeated and subjects rewarded for each choice at the end of every round. However, it is precisely these type of environments that exist when people function in markets and make choices repeatedly which are reinforced by an immediate payoff. Hence, we feel that our results have direct relevance to learning in market environments and the beneficial aspect of social learning or advice giving in them. (Stock brokers may not be that bad after all.)

The reason why we feel that advice giving and receiving facilitates learning is that the process of giving advice seems to focus the attention of advisers on the problem at hand in a manner that leads to greater learning on their part. Subjects seem to learn better when they give advice and when they receive it. We think this is true because giving and accepting advice causes decision makers to not only think through the problem another time but to do so in a manner different from the way they do when they are making decisions alone.

This result offers a possible explanation of why it is advantageous to follow advice when it is offered and why advice is better than actions. The reason is that we can expect subjects to learn better when they give advice and that advice is therefore worth listening to. The process of advice giving makes us think about the problems facing us differently than we tend to do when we are actually engaged in them.

6. WEAK-LINK GAMES WITH ALMOST COMMON KNOWLEDGE

All of the above experiments are ones where subjects receive private advice from one and only one predecessor. However, in many situations we get advice from several or many people and this advice is many times public. To investigate the role of advice in these situations, Chaudhuri, Schotter, and Sopher (2002) study what Van Huyck, Battalio, and Beil (VBB) (1990) have called "the Minimum Game" whose payoff structure is presented in Table 8.

This game is generated by a situation in which a set of agents choose an integer from the set $e_i \in \{0, 7\}$, $i = 1, 2, \ldots, n$. The payoff to each agent, i, is equal to $\pi_i = a + b\{\min(e_1, \ldots, e_n)\} - c(e_i)$, where $a, b, c > 0$ are constants. If $a = \$0.60$, $b = \$0.20$ and $c = \$0.10$, we get the matrix defined above.

In this game all outcomes in which all agents make the same choice are equilibria but the best equilibrium outcome for society is where all choose 7 while the worst is where all choose 1.

When Van Huyck et al. ran this game with groups of 14 or 16 they all quickly converged to the worst all-1 outcome and this result is remarkably robust. Groups tend to choose the worst outcome for themselves.

Now think of this game being played as an inter-generational game with advice. Here, one might expect that if agents in generations could talk to each other they might, even when their generation has failed to make the correct decision, give advice to

Table 8. Payoff Table in VBB's Minimum Game

		Minimum Choice of Others						
		7	**6**	**5**	**4**	**3**	**2**	**1**
Your Choice	**7**	1.30	1.10	0.90	0.70	0.50	0.30	0.10
	6	–	1.20	1.00	0.80	0.60	0.40	0.20
	5	–	–	1.10	0.90	0.70	0.50	0.30
	4	–	–	–	1.00	0.80	0.60	0.40
	3	–	–	–	–	0.90	0.70	0.50
	2	–	–	–	–	–	0.80	0.60
	1	–	–	–	–	–	–	0.70

the next generation that instructs them to learn from their mistakes and choose higher. That is, we expected that subjects would tell their laboratory offspring, "Do as we say but not as we did". This intergenerational talk, we expected might be able to allow subjects to "talk themselves to efficiency" and achieve optimal outcomes.

6.1. Did Advice Improve Welfare?

Chadhuri, Schotter, and Sopher (2002) find that it was much harder for societies to "talk themselves to efficiency" than they expected. More precisely, they find that efficient (all choose 7) outcomes emerge only in circumstances where advice is not only public (in the sense that all advice from a previous generation is offered to each successor in the next generation) but its publicness is common knowledge (in the sense that it is read aloud for all members of a generation to hear). Private advice between a predecessor and his or her successor or even public advice that is shared (i.e., all players in generation t are given a sheet specifying each piece of advice offered by the members of generation $t - 1$ and all subjects know that all others have been given the same sheet) but not read aloud does a poor job of raising the minimum.

Figure 6 exhibits the period-by-period minimum choices of subjects for five experimental runs, three of which are inter-generational games run with either advice only, advice plus last period's history (i.e., the ability to see what happened in each of the 10 rounds of the previous generations experience) and public advice. (There are 8 subjects in each group). The other two are the VBB (1990) results and the results of an experiment run to replicate them in our lab for comparison purposes.

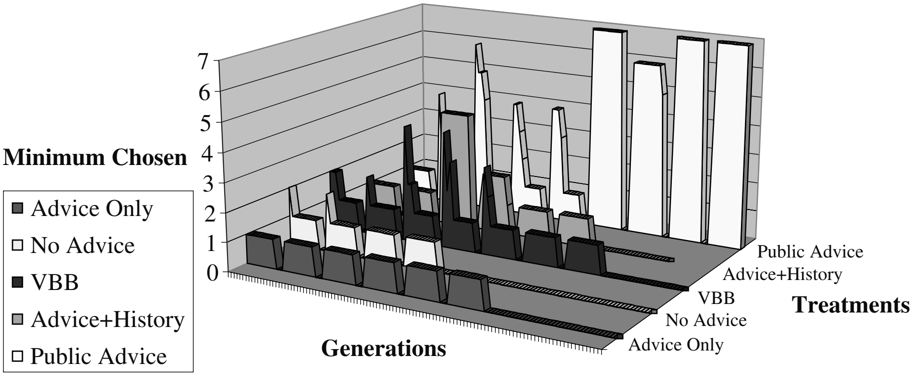

Figure 6. Behavior of the Minimum Across All Treatments.

These are called the replicator experiments and have neither advice nor history available.

As can be seen in Figure 6, the behavior of group minima in the VBB, our No-Advice, and all private advice experiments is dramatic. In all of them, the minimum converges to 1 by at least the fourth round. Our remarkable finding is that efficient all-7 outcomes become possible only when people have access to public advice that is also common knowledge. The discontinuity of behavior as we move from Almost-Common Knowledge (generations 1–5 of the Public-Advice Treatment) to Common Knowledge (generations 6–9 of the Public-Advice treatment) was unexpected but extremely suggestive of what type of information may be needed to get people to cooperate.

These experiments are extremely suggestive, implying that when we want groups of people to act in a coordinated manner in situations where they need to trust each other, the form in which the advice given is crucial. Private advice tends to make things worse – people advise their offspring not to be suckers and to watch out for numero uno. Public advice can help but only when it is common knowledge and all have faith that each other has not only heard the public advice but knows that others have heard its and that others know that they know that it was heard ad infinitum.

7. CONCLUSION

This paper has surveyed a number of papers all of which have investigated the impact of advice on decision-making. In general, this advice is offered by decision makers who are only slightly more experienced in the task at hand than are the people they advise. We call such advice "naïve". Despite this lack of expertise we find a number of results.

First people tend to follow the advice offered to them. This is seen in a number of ways. For example, in Ultimatum and Trust Games the amounts of money sent from the senders to the receivers is remarkably close to the amounts these subjects are advised to send. In coordination games, subjects tend to choose the action they are told to even when that advised action differs from that action that constitutes a best response to their beliefs about their opponent (which we elicit).

Second, not only is advice listened to and followed but it tends to change people's behavior. Games played with advice are played differently than those same games played without it. In addition, efficiency is generally higher when games are played with advice. This is true in coordination games where subjects tend to coordinate more often, in social learning tasks where advice increases the incidence of herds and cascades but always on the right decision, in one-person learning tasks, and finally in the Minimum Games when advice is public.

Finally, we feel that the reason why advice is so beneficial is that it forces decision makers to look at the problem they are facing in a more detached manner. The act of giving advice forces one to rethink the problem at hand while the act of receiving advice forces one to evaluate the advice that is given. Both endeavors lead a decision maker to take a more global approach to the problem and help a decision maker see the forest for the trees.

There are many things left undone in this research program. For example, we have ignored the process of advisor expertise. Many times the decision of whether to follow advice or not depends on the reliability of the person giving it. In our experiments this quality is suppressed. However, we do point out that even bad advice can be beneficial if it forces the person receiving it to sit down and think of the reason why it should be ignored. Second, there are many reasons to follow advice that have nothing to do with its informational benefits. For example, following advice can absolve a decision maker of the blame that results when decisions turn out badly. Many times the rationale for hiring experts is not that their opinion is so valuable but rather the act of hiring an expert can be used by a business executive to justify the process upon which his or her decision was made. These extensions and many more will have to be postponed for future research.

NOTES

[1] We use a non-overlapping generation structure and not an overlapping generations one because in most overlapping generation games of this type (see Salant (1991), Kandori (1989), Cremer (1986)) cooperation is achieved by each generation realizing that they must be nice to their elders since they will be old one day and if the current young see them acting improperly toward their elders, they will not provide for them in their old age. The analysis is backward looking in that each generation cares about the generation coming up behind them and acts properly now knowing that they are being observed and will inter-act directly with that generation. In this literature, folk-like theorems are proven if the length of the overlap between generations is long enough. In our work, however, generations never overlap. What they do is hope to behave correctly so that their children will see them as an example and act appropriately toward each other. Since they care about their children, adjacent generations are linked via their utility functions but not directly through strategic interaction. Hence, our model is a limiting type of overlapping generations model where the overlap is either minimal or non-existent.

[2] Except for the use of advice and the inter-dependence of our generational payoffs, our game shares many of the features of the procurement games of Jackson and Kalai (1997).

[3]

Regression of Amount Sent on Advice Dummy			
	Coeff.	*t-test*	*Prob. > t*
Dependent Variable: Amt. Sent			
Dummy (1 if advice allowed)	−13.18	−2.98	.00
Constant	40.18	10.82	.00
Model *F* Stat (Prob > *F*)	8.90 (.00)		
Adj. R-Squared	.03		
N	225		

[4] Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992) introduced the basic concepts and stimulated further research in this area. For surveys see, Gale (1996), and Bikhchandani, Hirshleifer and Welch (1998). Chamley (2003) discusses the different models of social learning and the relations between the different results.

REFERENCES

Anderson L. and C. Holt, (December 1997). "Information Cascades in the Laboratory." *American Economic Review*, 87(5), 847–62.

Banerjee A., (August 1992). "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, 107(3), 797–817.

Bikhchandani S., D. Hirshleifer and I. Welch, (October 1992). "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascade." *Journal of Political Economy*, 100(5), 992–1026.

Bikhchandani S., D. Hirshleifer and I. Welch, (Summer 1998). "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades." *Journal of Economic Perspective*, 12(3), 151–70.

Çelen B., Kariv S. and Schotter, A. (2003). "The Advice Puzzle: An Experimental Study of Social Learning Where Words Speak Louder Than Actions." *Mimeo Center for Experimental Social Science*, New York University.

Çelen B. and S. Kariv, (2003). "Distinguishing Informational Cascades from Herd Behavior in the Laboratory." *American Economic Review*, forthcoming.

Chamley C., (November 2003). Rational Herds Economic Models of Social Learning. Cambridge University Press.

Chaudhuri, Ananish, Schotter, Andrew, and Sopher, Barry, (2002). "Talking Ourselves to Efficiency: Co-ordination in Inter-Generational Minimum Games with Private, Almost Common and Common Knowledge of Advice." *Mimeo Center for Experimental Social Science*, New York University.

Cooper, Russell, DeJong, Douglas V., Forsythe, Robert, and Ross, Thomas W. (Winter 1989). "Communication in the Battle of the Sexes Game: Some Experimental Results." *Rand Journal of Economics* 20, 568–587.

Cooper, Russell, DeJong, Douglas V., Forsythe, Robert, and Ross, Thomas W. (May 1992). "Communication in Coordination Games." *Quarterly Journal of Economics* 107, 738–771.

Cremer, Jacques. (February 1986). "Cooperation in Ongoing Organizations." *Quarterly Journal of Economics*, 101: 33–49.

Gale D., (April 1996). "What Have We Learned from Social Learning?" *European Economic Review*, 40(3–5), 617–28.

Iyengar, Raghuram and Andrew Schotter, (2002). "Learning With A Meddlesome Boss: An Experiment in Sloppy Principal-Agent Relationships." *Mimeo Center for Experimental Social Science*, New York University.

Jackson, Matthew O. and Kalai, Ehud. (October/November 1997). "Social Learning in Recurring Games." *Games and Economic Behavior*, 21, 102–134.

Kandori, Michihiro. (January 1992). "Repeated Games Played by Overlapping Generations of Players." *Review of Economic Studies*, 59, 81–92.

Kandori, Michihiro, Mailath, George J., and Rob, Rafael. (January 1993). "Learning, Mutation, and Long Run Equilibria in Games." *Econometrica*, 61, 29–56.

Merlo, A. and Schotter, A. (1999). "A Surprise-Quiz View of Learning in Economic Experiments." *Games and Economic Behavior*, 28, 25–54.

Merlo, Antonio and Schotter, A. (2003). "Learning By Not Doing: An Experimental Investigation of Observational Learning." *Games and Economic Behavior*, forthcoming.

Salant, David. (1988). "A Repeated Game with Finitely Overlapping Generations of Players." *Games and Economic Behavior*.

Schotter, Andrew, and Sopher Barry. (June 2003). "Social Learning and Coordination Conventions in Inter-generational Games: An Experimental Study." *Journal of Political Economy*.

Schotter, Andrew and Sopher, Barry, (2004a). Advice and Behavior in Intergenerational Ultimatum Games: An Experimental Approach. Mimeo, *Center for Experimental Social Science*, New York University.

Schotter, Andrew and Sopher, Barry, (2004b). Trust and Trustworthiness in Games: An Experimental Study of Intergenerational Advice, Mimeo, *Center for Experimental Social Science*, New York University.

Van Huyck, John B., Battalio, Raymond C., and Beil, Richard O. (March 1990). "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure." *American Economic Review*, 80, 234–248.

Chapter 13

FAILURE OF BAYESIAN UPDATING IN REPEATED BILATERAL BARGAINING

Ching Chyi Lee
The Chinese University of Hong Kong

Eythan Weg
Purdue University

Rami Zwick
The Hong Kong University of Science and Technology

Abstract

Ever since Camerer and Weigelt (1988) concluded in their important experimental work that "sequential equilibrium describes actual behavior well enough," we might be tempted to use this theory confidently in various domains. To assess the robustness of the above conclusion, the present study attempts to explore Bayesian updating in a bilateral negotiated sale setup injected with a whiff of an ultimatum aroma. We conclude that the ultimatum nature of the basic game tends to overwhelm rational behavior on the part of the sellers and that buyers are not cognizant of favorable prices occurring later in the game.

1. INTRODUCTION

Ever since Camerer and Weigelt (1988) concluded in their important experimental work that "sequential equilibrium describes actual behavior well enough," we might be tempted to use this theory confidently in various domains. The present study explores Bayesian updating in a bilateral negotiated sale setup injected with a whiff of an ultimatum aroma, in order to assess the robustness of the above conclusion.

It is quite natural for people or institutions to misrepresent their true nature in pursuit of gaining some benefits which otherwise could not be attained. To misrepresent one's true nature is to act as someone or something else – thereby creating confusion on the true identity of the actor.

Situations in which a party may have an incentive to misrepresent are prevalent: business-to-business suppliers supply in the present and expect to be paid in the future; credit card issuers rely on credit card holders to pay for their purchases.

A. Rapoport and R. Zwick (eds.), Experimental Business Research, Vol. II, 249–260.

These are just two examples. Whether a business or an individual is likely to misrepresent depends on the situation and, perhaps, on the law.

Although misrepresentation may touch on questions of the law (for example, the energy company Enron, that borrowed huge sums of money based on a complicated web of holding companies although the chances of the company paying back the loans were very small), there are situations in which misrepresentation may only be a matter of benign convenience and opportunity as the framework explored in the present paper will clearly show.

Bargaining appears to be a natural scenario for misrepresentation because *reputation* is important. Feigning toughness for a sufficient length of time may convince the opponent in the bargaining to relent because of the cost incurred. Raiffa (1982), in his classic *The Art and Science of Negotiation*, writes that although repeated bargaining is often done cooperatively, this is not always the case, especially when there is information disparity between the two sides. "With repetition, a negotiator might want to establish a reputation for toughness that is designed for the long-term rather than short-term rewards" (p. 13). To be successful in this attempt, toughness needs to be communicated in some way to sow the seed in the opponent's mind that the toughness is real.

2. THE MODEL

The basic setting for this study includes a buyer and a seller. The buyer is one of two types: low cost (L) and high cost (H). In the current version of the experiment, this is operationalized as the low or high costs related to the seeking of an alternative supplier for an identical product that the seller proposes to sell. Upon receipt of the proposal to sell the product at a specific price, the buyer may accept it and thus terminate the transaction, or opt to *search* at a cost, c, for a better price by another supplier. The search for another supplier is always successful; however, the price may be better or worse than the current one proposed by the present seller. If the buyer elects to search, she abandons the opportunity to purchase the unit at the original seller asking price and is committed to pay the "searched" price even if it is higher than the current asking price (i.e., this is a no recall environment). A specific example is a bakery that can not roll over baked goods from one day to the next and the described encounter takes place just prior to closing. A buyer comes in and the baker is making a take-it-or-leave-it offer of price p for a cake. If the buyer accepts the offer the sale is conducted, otherwise, if the buyer rejects the offer the baker disposes of the cake and closes the shop. The buyer in this case is searching for the lowest available price elsewhere, incurring the search cost, and purchase the cake at that price. In general, the no recall environment allows for the negotiation to re-open if the buyer returns from an unsuccessful search. Presumably, the buyer's bargaining position in this case is much weakened. The seller is characterized by uncertainty about the real nature of the buyer. Obviously, a buyer with a known low cost would extract a better price than one characterized by a high search cost. Hence, regardless of her type, a buyer always has an incentive to be known as a low-cost type. As a result, the buyer's "telling" of his type is useless because it cannot be *trusted*. The only way a low-cost search buyer can reveal her true identity is by behaving in a such a way that it is inconsistent with her being a high-cost search buyer. In our case, this can be achieved by the buyer's willingness to reject low price offers. In this way, the buyer behaves as if she has a low search cost and thus can eventually creates a *reputation* for being a low-cost buyer.

2.1. The Bargaining Game

There are two players: a buyer and a seller. The seller possesses five units of a product that he intends to sell to the buyer in five periods, one unit in each period. The buyer is known to the seller to be one of two *types*: H or L. Initially, the degree to which the seller believes that the buyer is of type L is $0 \leq \pi \leq 1$. The search cost in buying the product is 0 if the buyer buys the product from the current seller. Otherwise, if the buyer does not purchase the product from the current seller, she needs to search for it at a cost $c = h$ or $c = l$ $(h > l)$ corresponding to her type. A search means a realization of a random variable D having a uniform distribution on the interval [0, 100]. Thus, when the seller considers a price offer, p, he needs to assess the expected price and the cost that the buyer incurs in case the offer is rejected. The profit for the buyer is $100 - p$ when the price is accepted and $100 - D - c$ otherwise. Thus, the buyer has 500 pay units at his disposal and she wants to minimize the total price over the five periods of the game with the same seller. The seller's profit from each period is the price he can obtain for the unit. A failure to sell the product in a period results in zero profit for this period. The seller's goal is to maximize the total revenue over the five periods. All the information as described above is commonly known to the participants.

2.2. The Equilibrium

The equilibrium of the game described above is very similar to that of the game described by Kreps and Wilson (1982) and also to that of the one described by Camerer and Weigelt (1988). We will only describe the equilibrium path here. The details and the derivation of the equilibrium can be found in Lee (2000).

Let p_H and p_L be the prices that make the high- and low-type buyers indifferent between searching and not searching. That is, $p_H = -h + [100 - E(D)]$ and $p_L = -l + [100 - E(D)]$, where $E(D)$ is the expected price if the buyer searches. Let π_t be the seller's belief at Period t about the probability that the buyer is a "low-cost buyer". Notice that we count the "periods" backwards. That is, in period t, there are t periods left. Furthermore, let $\theta^t = [(p_H - p_L)/p_H]^t$. Then, in equilibrium, the seller offers p_L until the first time that π_t becomes lower than θ^t, at which time the seller offers p_H until the last period. The low-cost buyer accepts $p_t = p_L$ and rejects $p_t = p_H$. The high-cost buyer accepts $p_t = p_L$ with probability one and rejects p_H with probability β_t if she has not yet accepted any price greater than p_L. If the high-cost buyer has accepted a price greater than p_L earlier, she accepts p_H with probability one. Here, β_t satisfies:

$$\beta_t = \frac{\pi_t(1 - \theta^{t-1})}{(1 - \pi_t)\theta^{t-1}}.$$

3. THE EXPERIMENT

3.1. Method

3.1.1. Subjects

Two hundred and forty male and female subjects, who were mainly undergraduate business students at the Hong Kong University of Science and Technology, in groups of 24 students per session, participated in a session that lasted about 60 minutes. Subjects were recruited through advertisements placed on bulletin boards on campus and made during class announcements. The announcements promised monetary reward contingent on performance in a bargaining study.

3.1.2. Experimental Design

Each of the bargaining games consisted of five periods in which the same players bargained on a surplus of HK$100[1] in each period with an uncertain outside option for the buyer uniformly distributed on the range [0, 100] using the trading rules described above.

We used a 2 (High search cost) × 5 (Degree of belief) × 8 (trials) design. The first two factors were between subjects and the last was within subjects. In all sessions the level of the low search cost was fixed at $5 (a commonly known fact). The high search cost was either $10 or $45. The prior belief about the buyer being a low-cost type was 0.00, 0.01, 0.05, 0.10, or 0.50.

Subjects assumed the same role (a buyer or a seller) in all eight trials in a session and faced different anonymous opponents on each trial. Obviously, the five periods of each trial were played by the same buyer and seller.

During each trial, the seller was asked to sell five indivisible goods, one in each period. The goods had no value to the seller except their selling prices. The value of the good to the buyer was $100. The bargainers knew both reservation prices. Before the beginning of each trial the computer randomly selected the buyer's type based on the known probabilities of each type. Of course, the buyer was informed of her type, whereas the seller only knew the sampling probabilities. Once selected, the buyer's type was fixed for the five periods duration of the game. The game proceeded as follows: at the beginning of each period the seller announced a selling price for the good (the asking price). The buyer then had the following options:

1. Accept the asking price, thereby terminating the period.
2. Search for an alternative price. In this case, the buyer had to pay a search cost, and a price was randomly generated (the outside offer) from the range [0, 100]. The price generated through the search (if any) was known to both bargainers.

After learning about the outside offer, the buyer must accept this offer, thereby terminating the period.

At the end of each period, the buyer was informed of her profit as well as that of the seller. The seller was informed of his profit and that of the buyer's profit if the buyer accepted the seller's asking price, or the buyer's profit up to the uncertainty as to the buyer's search cost if the buyer decided to search (i.e., the buyer's profit was presented to the seller as "outside price – search cost." Whereas outside price was explicitly specified as a number, the search cost was understood to be the unknown buyer's type and as such was literally presented as "search cost"). After the profits were presented the game proceed to the next period, unless it was the last (fifth) period, whereas the game was terminated.

The subjects interacted in a computer laboratory arranged in such a way that it was impossible for the subjects to know with whom they were negotiating or to see each other's screens. Asking prices, acceptances, and searches (including their outcomes) were transmitted through computer terminals. No other communications were allowed.

Throughout the experiment, the subjects were informed of every *known* characteristic of the game being played. Moreover, the *known* dynamic aspects of the settled and searched prices (during the five periods of each game) were registered in a history log and were clearly visible on the screen for the duration of the game.

The subjects were informed that they would be paid their net payoff from one randomly selected game (cumulative over the five periods). In addition, each subject was paid $10 for participation. On the average, subjects earned $73.64 for a session[2].

4. RESULTS

4.1. Data Consistency

The data includes 4,800 plays of the same basic game arranged in 10 sessions, covering 960 plays of the repeated game.

Subjects knew that searching will result in a randomly drawn price from the interval [0, 100] and that within this interval every price is "equally likely". That is, the expected search price was 50. If we add the commonly known minimal cost to pursue the search (5) or the maximum possible cost to search (45), we arrive at the fundamental conclusion, which requires no deep thinking or understanding, that asking prices (demands) should be no less than 55 and no more than 95.

Inspecting the data for coherence, we find two anomalies, one minor and one major:

- There are two unusual price demands of 553 and 47,055. These are beyond the range of the upper bound of the random search price, and are also beyond any reasonable demand. The seller's computer interface allows entering a price expressed in dollars and cents, but it was the seller's task to enter the decimal point. However, the interface failed to alert sellers to unusual demands such as we have here. We believe it reasonable to assume that subjects miss writing a decimal point and intended to demand 47.55, and 55.3, respectively.

- A total of 3006 demands (close to two-thirds of all the plays) were *below* the lowest reasonable price demand of 55, 1673 of them below 50.

Table 1 presents the number of price demands below several cutoff points:

Table 1. Distribution of prices below the lower bound price demands

Cutoff	< 50	< 45	< 40	< 35	< 30
Count	1,673	882	448	212	75

Table 2 presents the frequencies of demands below the minimum expected demand of 55 in the ten experimental conditions.

The two distributions presented in Table 2 shed light on the problem. These demands, although irrational, are consistent (column-wise) with the fact that the frequency of price demands below 55 is lower when h is 45 than when it is only 10. Thus, when the real higher bound is likely to be higher than 60, subjects tend to deviate from the lower bound of 55 less often. Or, in other words, when the real higher bound is likely to be higher, subjects tend to make higher price demands more often.

Table 2. Frequency of demands below 55

High Search Cost (H)

π	*45*	*10*
0.00	185	433
0.01	143	320
0.05	273	434
0.10	168	399
0.50	228	423

4.2. Typical Demands

Given the high level of irrational behavior, which is further accentuated when one takes into consideration mild risk aversion, there is little surprise with the following results.

We next look at the mean behavior. Recall that in each experimental condition (i.e., crossing the seller's five levels of initial belief about the buyer's type with the two levels of the high search cost) consists of 12 pairs of subjects in fixed roles (buyer or seller) playing eight bargaining games. Each bargaining game consists of five periods in which a fixed pair plays the price-taking game as set by the seller. The buyer may rebuff the seller by irrevocably opting for an outside alternative through searching (at a cost).

In order to stabilize the data, we consider the median price demand of any given seller over the eight repetitions of the plays he participated in during the session and in each period (1–5). This eliminates the possible effects of extremely small or extremely high demands.

Table 3 presents the means of these medians over the 12 sellers in each session, by initial sellers' beliefs (π), high search cost (H), and period.

For comparison, Table 4 presents the equilibrium price demands, assuming risk neutrality.

By observing the equilibrium price demands in Table 4, we notice the following:

1. For $\pi = 0.00$, the price demand in each period should be higher when $H = 45$ than when $H = 10$.
2. For $\pi = 0.50$, the price demand in each period should be the same no matter whether $H = 45$ or $H = 10$.
3. For $\pi = 0.01$, 0.05, and 0.10, the price demands are generally lower in the earlier periods than in the later periods (they go up from 55 to either 60 or 95).

Table 3. Mean medians price demands

$\pi\rightarrow$	*0.00*		*0.01*		*0.05*		*0.10*		*0.50*	
$H\rightarrow$ *Period*	*45*	*10*	*45*	*10*	*45*	*10*	*45*	*10*	*45*	*10*
1	56.8	46.8	61.1	52.6	51.2	43.2	56.0	46.6	53.6	49.1
2	56.5	45.8	59.9	49.9	51.8	43.2	56.6	46.0	54.7	49.8
3	56.2	45.9	59.6	49.5	53.2	43.8	56.5	45.2	55.4	49.5
4	55.8	47.2	60.7	51.2	53.4	42.7	56.7	46.5	55.3	49.6
5	56.0	47.3	60.0	50.8	53.5	42.2	56.5	45.6	55.0	49.4

Table 4. Equilibrium price demands

$\pi\rightarrow$	0.00		0.01		0.05		0.10		0.50	
$H\rightarrow$ Period	45	10	45	10	45	10	45	10	45	10
1	95	60	95	55	55	55	55	55	55	55
2	95	60	95	55	55	55	55	55	55	55
3	95	60	95	55	95	55	55	55	55	55
4	95	60	95	55	95	55	95	55	55	55
5	95	60	95	60	95	60	95	55	55	55

The most striking observation from the data presented in Table 3 is the stability of the means over the five periods. The subjects do not learn! As mentioned in point 3 above, one would have expected some change as the seller learns about the buyer. We do not show it here, but this stability is also noted at the individual levels.

In Table 3, the first two columns on the left describe behavior under the well-known ultimatum conditions (Güth et al., 1982). The results deviate markedly both from rational expectation on the one hand and from traditional ultimatum results on the other hand, in which one expects about 2/3 of the range to be given to the seller. This translates to price of about 97 and 86.

The rest of the columns show that, except for one case (the Period 1 demand of 61.1 when $\pi = 0.01$ and $H = 45$), all means are outside the rational interval. This again shows that rational Bayesian behavior is far too much to expect. It seems that subjects frame the situation they are in quite differently.

The only understandable aspect of Table 3 is the relation between corresponding columns members under equal π. For a given π, the higher the search cost, the higher the mean price demand. However, this is a very weak prediction. It is also compatible with the corresponding number of deviations from the rational interval of prices, as we have seen above.

4.3. Behavior Change

We next look at the relations between price demands and acceptance behavior. Consider the correlation between the search behavior and the price demand. We code the search as 1 if the buyer searches and 0 otherwise. Table 5 presents the correlations between search behavior and price demand by the conditions of the experimental design, except that we collapse the data from periods 2 to 5. The goal

Table 5. Correlations between search behavior and price demands

$\pi\rightarrow$	*0.00*		*0.01*		*0.05*		*0.10*		*0.50*	
$H\rightarrow$ *Period*	*45*	*10*	*45*	*10*	*45*	*10*	*45*	*10*	*45*	*10*
2–5	0.300	0.297	0.286	0.187	0.271	0.235	0.156	0.364	0.203	0.087
1	0.442	0.413	0.504	0.510	0.330	0.457	0.440	0.563	0.474	0.364

Table 6. Search counts

$\pi\rightarrow$	*0.00*		*0.01*		*0.05*		*0.10*		*0.50*	
$H\rightarrow$ *Period*	*45*	*10*	*45*	*10*	*45*	*10*	*45*	*10*	*45*	*10*
1	35	34	38	47	28	45	26	41	39	18
2	26	26	27	37	24	40	18	30	34	23
3	25	31	25	35	27	35	22	37	38	24
4	25	31	24	39	26	34	20	30	33	28
5	11	30	25	37	15	30	13	41	34	24

is to observe if there is any change in the response behavior between the beginning and the rest of the game.

We do not test any formal significance. It is quite noticeable that (a relatively) high price demand is met with a tendency for the buyer to decide to search since all correlations are positive. This is, of course, not an indication of formal updating, since, as we have seen, most demands should have been accepted anyway.

Table 6 shows the total number of searches in each period by π and H. The base is always 96 games played in each cell. We see that buyers initially tend to reject and search. This seems less natural when the buyer is equally likely to be weak or strong (the last two columns when $\pi = 0.50$).

5. DISCUSSION

We divide our discussion into two parts: the sources of difficulties in playing the game "correctly" and some modeling issues. The latter arises from the observation that sellers violated the most basic prediction of the rational model, even when there was no uncertainty as to the buyer's type ($\pi = 0$). Since sellers' asking prices were so low vis-à-vis the rational prices, buyers were "deprived" of the opportunity to strategically reject prices in early rounds in order to manage the seller's beliefs about the their type. The low asking prices were akin to the seller stating that he believes that the buyer's search cost is low, even though the objective probability of a low-cost type was much less than 50% in most cases.

5.1. The sources of difficulties in playing the game "correctly"

We have seen that even superficial analysis of the bargaining game allows us to conclude that there is only meager evidence to indicate some support for rational behavior in this setting (see the correlation data above). Of course, what is rational is not absolutely clear. Let us check our assumptions again. The bargaining interval, which is the interval in which price demands must fall in order to satisfy the minimum price that the seller may demand and the maximum price beyond which the buyer will not settle, was derived from the distribution of the search price. The derivation assumes that it is common knowledge that all participants are risk neutral. The existence of a commonly known bargaining interval assumes that the participants would rather gain more than less out of playing the game.

Now we expressly admit that there is no reason in terms of game theory to require the players to figure out the bargaining interval in such a complex manner. We know that an average person's understanding of probability has its limitations, which may have nothing to do with game behavior. In our setting, in addition to the strategic uncertainty and the uncertainty (from the seller's point of view) as to the buyer's type, a third layer of uncertainty existed as to the realization of the actual outside price. We now believe that this added complexity might have unnecessarily clouded the strategic nature of the game. In fact, a next iteration of this experimental program intends to simplify the bargaining game by eliminating the probabilistic nature of the outside option. Consider, for example, a game similar to the one we have implemented except that the unknown (to the seller) buyer's type is characterized by her outside option. Now the bargaining interval is easy to figure out! Any solution of the original game is a solution of this proposed one. Eliminating one source of uncertainty might re-focus subjects' attention on the reputation nature of the scenario.

Another problem is hinted at by the extreme case when π is 0. In this case, the game is simply an ultimatum game with an uncertain outside option to the responder. This game was investigated by Zwick and Lee (1999) in their No Recall condition. They report that the classical game-theoretical model can account for some (but not all) of the behavioral regularities. In line with recent developments in behavioral decision theory and game theory, which assume bounded rationality and preferences over the relative division of a surplus, Zwick and Lee (1999) found that subjects followed simple rules of thumb and distributional norms in choosing strategies, which are reflected in the behavioral patterns they observed. We know that ultimatum games are notoriously difficult to play rationally. Some of the other π values in our design may not be sufficiently different from 0. So, the framing of the bargaining situation with an ultimatum skeleton obviously, in hindsight, makes our data difficult to interpret.

It is not clear how to re-frame the game to avoid the ultimatum aspect. One way is to change the number of periods played and the information given to the seller.

5.2. *The Modeling Issues*

Needless to say, the rational reasoning discussed in the introduction is not sufficiently powerful or intuitive to derive behavior in practice. In fact, the game we described requires some elaboration and additional scaffolding to obtain more natural predictions. To see why the data are not sufficient, we note that if we assume that the buyer may only choose to accept or search by some deterministic rule, we know that this rule must obey the property that if she accepts any price, then she must accept any lower price. And since the buyer is one of two types, once she accepts a price not on the boundary of the bargaining interval, she reliably signals her identity as a high-cost buyer. Since a deterministic rule is not sufficient to uncover what seems to be natural behavior, a probabilistic rule is one approach to enrich buyer's behavior.

One such operationalization is given by Lee (2000), except that in his approach both the buyer and the seller are equipped with probabilistic strategies. In any approach that relies on such rules, both players need to be commonly cognizant of the extension to the given game. So, aside from the computational complexity, the players need some common understanding of the extension. But are these "computational complexities"? In the 1960s, in a prelude to Kahneman and Tversky's work that initiated the heuristics and biases approach to decision theory, it was discovered that quite consistently humans are notoriously conservative in updating their degrees of belief (Edwards, 1968). Now, the work of Lee, as well as others, that relies on mixed extensions of sequential games also relies explicitly on Bayesian updating. While the experimental work discussed by Edwards (1968) explicitly described the probabilistic scenario in which people fail, the modeling in this case is implicit and the players need to commonly agree upon their existence and then update.

We conjecture that, in a plainer environment, one can at best expect that people are quite a bit slow in achieving the reputation they desire. But at least this may be experimentally attainable.

ACKNOWLEDGMENT

The author acknowledges the financial support received from the Hong Kong Research Grants Council for this research (Project Number: CUHK4076/98H).

NOTES

[1] \$100 Hong Kong dollars. The exchange rate between Hong Kong and US dollars is approximately 7.8 to 1.

[2] The amount of the cash prize was very attractive to students considering that the hourly wage for an on-campus job was about \$50.

REFERENCES

Camerer, Colin and K. Weigelt, (1988). “Experimental Tests of a Sequential Equilibrium Reputation Model.” *Econometrica*, 56(1), 1–36.

Edwards, W., (1968). “Conservatism in Human Information Processing,” in *Formal Representation of Human Judgment*, B. Kleinmuntz ed., John Wiley, New York.

Güth, W., R. Schmittberger, and B. Schwarze, (1982). “An Experimental Analysis of Ultimatum Bargaining.” *Journal of Economic Behavior & Organization*, 3(4), 3867–388.

Kleinmuntz, B., (1968). *Formal Representation of Human Judgment*, John Wiley, New York.

Kreps, D., and R. Wilson (1982). “Reputation and imperfect information.” *Journal of Economic Theory*, 27, 253–279.

Lee, Ching Chyi, (2000). “Reputation in Repeated Bilateral Trading with Outside Alternatives.” *Working Paper*, The Chinese University of HK.

Raiffa, H., (1982). *The Art and Science of Negotiation.* Harvard University Press, Cambridge MA.

Zwick, Rami and Ching Chyi Lee, (1999). “Bargaining and search: An experimental study.” *Group Decision and Negotiation*, 8(6), 463–487.

Author Index

Subject Index

The Authors

The Editors

Amnon Rapoport
Department of Management and Policy
University of Arizona
Tucson, AZ 85721
and
Department of Economics
Hong Kong University of Science and Technology
Clear Water Bay, Kowloon
Hong Kong
amnon@u.arizona.edu

Rami Zwick
Department of Marketing
Hong Kong University of Science and Technology
Clear Water Bay, Kowloon
Hong Kong
mkzwick@ust.hk

Chapter Contributors

J. Neil Bearden
Department of Management and Policy
The University of Arizona
Tucson, Arizona 85721
jneilb@gmail.com

Paul Brewer
Department of Economics
Hong Kong University of Science and Technology
Clear Water Bay, Kowloon
Hong Kong
pjbrewer@ust.hk

Kay-Yut Chen
Decision Technology Dept
MS1U-2, Hewlett-Packard Laboratories
1501 Page Mill Road
Palo Alto, Ca 94304
kychen@hpl.hp.com

Yan Chen
School of Information
The University of Michigan
1075 Beal Avenue
Ann Arbor, MI 48109-2112
yanchen@umich.edu

James C. Cox
Department of Economics and Economic Science Laboratory
Eller College of Management
University of Arizona
Tucson, AZ 85721-0108
jcox@eller.arizona.edu

Rachel Croson
Department of Operations and Information Management
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6340
crosonr@wharton.upenn.edu

Daniel Friedman
Economics Department
University of California
Santa Cruz CA
dan@ucsc.edu

R. Lynn Hannan
J. Mack Robinson College of Business
Georgia State University
Atlanta, GA 31709
accrlh@langate.gsu.edu

Stephen C. Hayne
Computer Information Systems
College of Business
Colorado State University
Fort Collins, CO 80523
hayne@acm.org

Vicky B. Hoffman
Katz Graduate School of Business
University of Pittsburgh
Pittsburgh, PA 15260
vickyh@katz.pitt.edu

Daniel Houser
ICES (Interdisciplinary Center for Economic Science)
George Mason University
4400 University Drive, MSN 1B2
Fairfax, VA, 22030
dhouser@gmu.edu

Suzhou Huang
2101 Village Road, MD 2122
Research and Advanced Engineering
Ford Motor Company
Dearborn, MI 48121
shuang10@ford.com

Bernardo Huberman
Hewlett-Packard Laboratories
1501 Page Mill Road
Palo Alto CA
huberman@hpl.hp.com

Ching-chyi Lee
Dept of Decision Sciences and Managerial Economics
The Chinese University of Hong Kong
Shatin, NT, Hong Kong
cclee@baf.msmail.cuhk.edu.hk

Donald V. Moser
Katz Graduate School of Business
University of Pittsburgh
Pittsburgh, PA 15260
dmoser@katz.pitt.edu

David Reiley
Department of Economics
University of Arizona
401 McClelland Hall
Tucson, AZ 85721
reiley@eller.arizona.edu

Tatsuyoshi Saijo
Institute of Social and Economic Research
Osaka University
6-1 Mihogaoka
Ibaraki, Osaka 567-0047
Japan
saijo@iser.osaka-u.ac.jp

Arie Schinnar
Associate Professor Emeritus of Business and Public Policy
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6372
schinnar@wharton.upenn.edu

Andrew Schotter
Department of Economics
New York University
269 Mercer Street
New York, New York 10003
andrew.schotter@nyu.edu

Darryl Seale
Department of Management
University of Nevada Las Vegas
4505 Maryland Parkway
Las Vegas, NV 89154-60009
dseale@unlv.nevada.edu

Eythan Weg
W. Lafayette, IN 47906
weg@indiscrete.org

John Wooders
Department of Economics
McClelland Hall
University of Arizona
Tucson, AZ 85721
jwooders@eller.arizona.edu