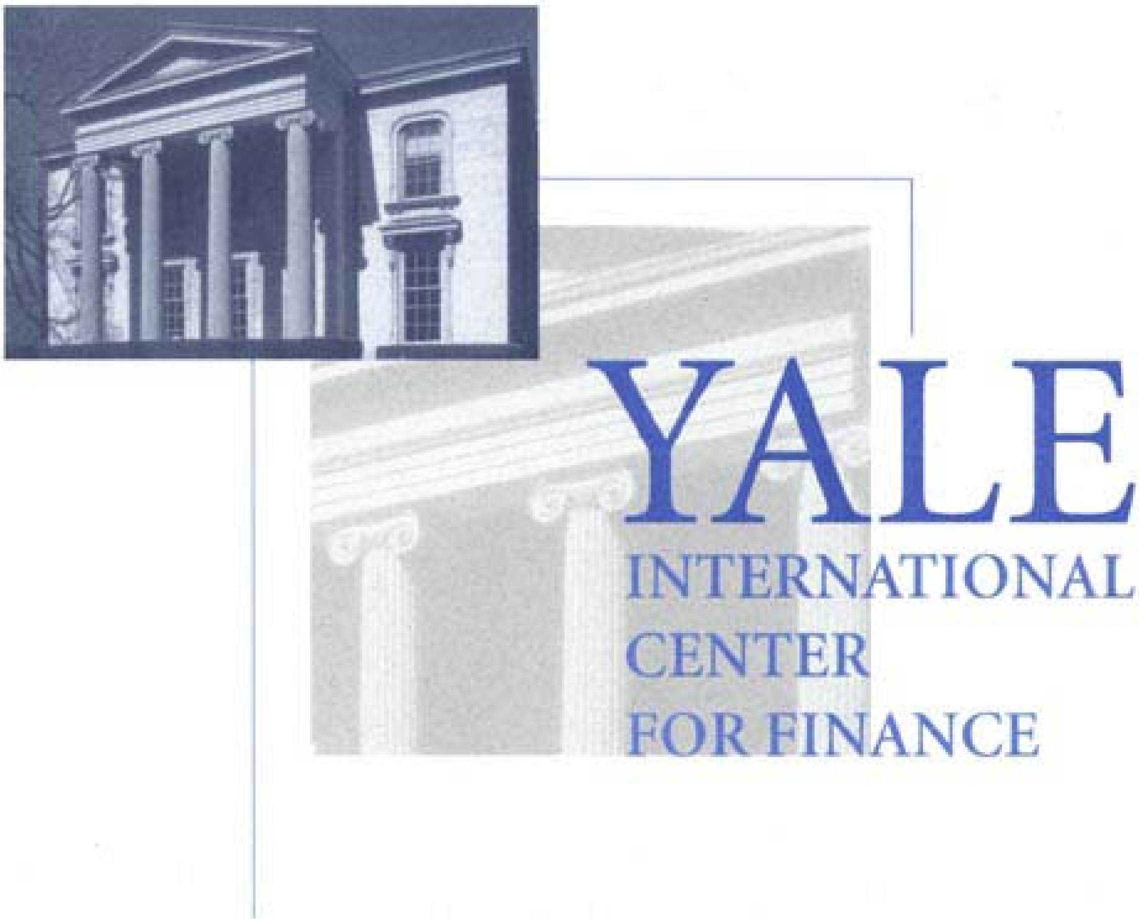

Yale ICF Working Paper No. 02-16
May 2003

MARKET AS ARTIFACTS: AGGREGATE EFFICIENCY FROM ZERO-INTELLIGENCE TRADERS

Shyam Sunder

Yale School of Management

Markets as Artifacts: Aggregate Efficiency from Zero-Intelligence Traders

Shyam Sunder
Yale University

Abstract

Three phenomena—the disparity between the assumed and observed attributes of economic man, the link between nature and artifacts, and the use of computers as a source of knowledge—fascinated Herbert A. Simon. He built a new paradigm for each field—bounded rationality to deal with the disparity, the science of the artificial as its link to nature, and artificial intelligence for creation of knowledge. In this paper we show that the sciences of the artificial and computer intelligence also hold a key to an understanding of the disparity between individual behavior and market outcomes. When seen as human artifacts, a science of markets need not be built from the science of individual behavior. We outline how, in the nineties, computer simulations enabled us to discover that allocative efficiency—a key characteristic of market outcomes—is largely independent of variations in individual behavior under classical conditions. The Sciences of the Artificial suggests such independence and points to its benefits.

JEL Classification: A12, B25, C9, D44, D81
Keywords: efficiency of markets, zero intelligence, decision making, bounded rationality, minimal rationality economics

Revised Draft
May, 2003

I am grateful to Philip Bromiley, Judy Carmel, Daniel Friedman, Dhananjay K. Gode, Karim Jamal, and Manjula Shyam for their help with this paper. This paper is available in PDF format from
http://www.som.yale.edu/faculty/sunder/research. Please send comments to
shyam.sunder@yale.edu.

Markets as Artifacts: Aggregate Efficiency from Zero-Intelligence Traders

Shyam Sunder
Yale University

> *…the possibility of building a mathematical theory of a system or of simulating that system does not depend on having an adequate microtheory of the natural laws that govern the system components. Such a microtheory might indeed be simply irrelevant.*
>
> Herbert A. Simon, The Sciences of the Artificial, p. 19.

Three phenomena—the disparity between the assumed and observed attributes of economic man, the link between nature and artifacts, and the use of computers as a source of knowledge—fascinated Herbert A. Simon. He built a new paradigm for each field—bounded rationality to deal with the disparity, the science of the artificial as its link to nature, and artificial intelligence for creation of knowledge. In this paper we show that the sciences of the artificial and computer intelligence also hold the key to an understanding of the disparity between individual behavior and market outcomes. When seen as human artifacts, a science of markets need not be built from the science of individual behavior. We outline how, in the nineties, computer simulations enabled us to discover that allocative efficiency—a key characteristic market outcomes—is largely independent of variations in individual behavior under classical conditions. *The Sciences of the Artificial* suggests such independence and points to its benefits:

> This skyhook-skyscraper construction of science from the roof down to the yet unconstructed foundations was possible because the behavior of the system at each level depended on only a very approximate, simplified, abstracted characterization of the system at the level next beneath. This is lucky, else the safety of bridges and airplanes might depend on the correctness of the “Eightfold Way” of looking at elementary particles (Simon 1996, p. 16).

Substantive and Procedural Rationality

In 1935, Simon faced the problem of understanding the allocation of the city budget between maintenance by the parks department and programs run by the public schools in Milwaukee. He could not see how the marginal benefits of two activities could be assessed, and how these incommensurables might be compared, much less equalized, according to the prescriptions of neoclassical economics (Larkey 2002). Economics assumes that agents choose the options they prefer most from their opportunity sets, and thus requires that they know the opportunity set at the time they choose. Simple algebra leads to the equalization of the marginal benefits as a logical consequence of this process.

Simon used "substantive rationality" as the label for such behavior. It is not clear how an individual is to achieve substantive rationality without knowing his opportunity set. What could the agent do when he knows but one option, and must search further—an economic decision in itself—to generate more options? Simon postulated that an agent starts out with an initial level of aspiration about his welfare and is willing to accept an option that satisfies him by attaining this level. Acceptance of an option concludes the search; rejection leads to lowered aspirations, search for another option, and application of the same stopping rule. Simon (1978) called this process "procedural rationality."

Field and laboratory observations support the descriptive validity of procedural over substantive rationality in human agents. They are rational in the sense of choosing what is best, but only boundedly so in the sense of choosing from a limited opportunity set in relation to their aspiration level. Yet economics routinely assumes that individuals choose from their opportunity sets to maximize their welfare, not merely to satisfy their aspirations. This is true not only of the elegant neoclassical foundations in general equilibrium theory, but also of applications in the theory of money, industrial organization, trade, labor, etc. Why build these theories from demonstrably false assumptions about agent behavior?

The positivist answer to such a question is: the descriptive validity of the model is not relevant as long as the model predicts well (Friedman 1953). Such answers are unsatisfactory because our models serve not only to predict but also to articulate our understanding of various phenomena, and to convey that understanding to others. Understanding of phenomenon is crucial to science; prediction without understanding does not build science. We show that the sciences of the artificial point to a better answer.

The Sciences of the Artificial

Artifacts comprise elements, each with its own inherent properties, governed by natural law. A boat has timber; a shirt has cotton or wool fiber; and a shoe has leather, along with other elements; which may be artifacts themselves. In natural sciences we analyze the elements of interest, in sciences of the artificial we synthesize artifacts from elements to attain goals or perform functions. In science, natural things simply are; it is not meaningful to ask how they ought to be. Of the artifacts, we can ask both how they are and how they ought to be.

Natural law governs the inner and the outer environments of artifacts, as well as the interface between the two (Figure 1). Presence of the goal or intent of its creator or user distinguishes an artifact from nature. How well an artifact fulfills these goals depends on the interface between the two environments.

A twig lying under a tree becomes an artifact when a chimpanzee picks it up and inserts it into a termite hill to extract food. Titanium alloy created to meet the performance demands of supersonic aircraft does not exist in nature. Both the twig and the alloy follow the laws of nature. The twig exists in nature; the titanium alloy is manufactured to meet the performance specifications of the aircraft. Both are artifacts to their creators and users.

It is possible to understand and predict the changes in the performance of an artifact as a function of the characteristics of its outer environment, contingent only on a few critical features of the inner environment.

The boundary between the inner and outer environments of an artifact is drawn by reference to the purpose behind its design, or its presumed function. If we were interested in all possible aspects of the relationship between the inner and the outer environments, the two would have a one-on-one correspondence. However, we are typically interested in only some limited aspects of this relationship for an artifact. This coarseness of interest creates redundancy in the correspondence: many inner environments may stand in a given functional relationship to a given outer environment, and many outer environments may also stand in a given functional relationship to an inner environment. The chimpanzee may use not only a twig but also a straw or a thin bone to extract termites from their hill. A twig may be used not only to get termites but also ants or honey from hard-to-reach spaces.

Important parts of the debate about the assumptions of economics are rooted in confusion about the roles of inner and outer environments of an artifact. For artifacts of physical substance, such as cars, cameras, or cities, the boundary between inner and outer environments is easy to see. For social artifacts without physical substance, the boundary is not so obvious. Consider markets as an example.

Markets as Artifacts

Markets are artifacts created by humans through social evolution or design. While both natural and artifactual phenomena are subject to the laws of nature, we can see all artifacts from a functional or teleological perspective.

Simon's characterization of artifacts suggests that in order to develop a science of markets and other such social systems it is useful to draw the boundary between their inner and outer environments. Market structure or rules lie inside, while the agents, defined by their endowments, preferences, and decision rules, lie on the outside. The usefulness of an artifact arises from its outcomes' standing in a desired relationship with the outer environment. The outcomes are determined by interactions between the inner and the outer environments under natural law. The choice of inner environment of the artifact generates the outcome function. The inner environment remains largely unnoticed by most users, usually attracting attention only when it is stretched beyond its limits and the outcome fails to stand in the desired relationship with the outer environment.

The rules of a market define its inner environment. These include a language consisting of admissible messages its participants can send, a mechanism to define and implement the distribution of these messages, a law of motion that defines which messages are valid in each state of the market, and a rule to allocate resources as a function of messages (Smith 1982).

In a supermarket for example, the seller sends messages about his willingness to sell through price labels. The buyer sends the messages about his willingness to buy at that price by transferring the appropriate quantities to his shopping cart and presenting it at the checkout counter. Price messages from the seller are made available to all buyers in the form of posted prices. Messages from the buyers are supposed to be available only to the checkout clerk, though it may be difficult to keep other buyers from looking at the cut of beef in the adjacent cart. The buyer cannot send a buy message when the store is out of stock. The allocation rule consists of payment of the sum of prices of groceries in the cart to the store and transfer of groceries to the customer. The inner structure of other markets, such as a stock exchange or bidding for construction of a municipal bridge, can be similarly defined by their rules.

Resource and information endowments, preferences, and decision rules of the participating agents form the outer environment of a market. The seller in a supermarket has information about estimated demand for each good at various prices. He chooses the goods, their prices and how they are displayed, using his decision rule to seek his goals, e.g., profits. Buyers combine the information about the prices and other relevant attributes of various goods with their tastes and budgets, and use their own rules to make buy decisions.

The interaction between these inner and outer environments of the market results in the transfer of money from various customers to the store's cash register or bank account, and the transfer of grocery baskets of varying composition to the customers. Prices and quantities of various items of grocery, the amount spent by each customer, the grocer's profit, and the net gain in satisfaction of the customer are some of the other outcomes of the market. We can assess a market as an artifact by examining how the outcomes of interest to us change as a function of the inner and the outer environment of the market.

We design a market by choosing its rules (inner environment) so a desired relationship between the outer environment and the selected outcomes is obtained. If the rules of the market arise from social evolution, we assess the functionality of this artifact on the basis of that relationship. Since the outcomes and the outer environment of any artifact are multidimensional, it is rarely possible or desirable to look at a complete mapping between them. We choose only a few critical features of the outcomes and outer environment to make the assessment. In designing a car seat for infants, for example, the safety of the child (the outcome) in car accidents of varying intensity (the outer environment) is an overriding consideration; matching the texture of the seat materials is not.

In neoclassical economics, the outer environment of a market is typically represented by the supply and demand conditions. Aspects of the outer environment not captured in supply and demand, such as the decision-making processes of the agents, are assumed to take simple and idealized forms. Allocative efficiency, price, and distribution of gains from trade are the prominent aspects of market outcome that receive attention in this tradition.

Many critical aspects of the outer environments of markets are unobservable in the field. The unique facility of computers in modeling the behavior of systems and their components and the use of the artificial intelligence paradigm helped identify which market outcomes are causally dependent on which attributes of their inner and outer environments.

New Knowledge from Simulation of Markets

Simon asked: *How can a simulation ever tell us anything that we do not already know?* It may help us compute the consequences of combinations and interactions among components of a system that may be too difficult to work out otherwise. In the case of markets, traders interact with other traders within the confines of the rules of the market. Even if the behavior of traders were well defined, their interactions can be quite complex, making it difficult to characterize the market outcomes in all except the simplest of market designs. Hence the theoretical prominence of the Walrasian auction, which is hardly seen in practice anywhere. Laboratory simulation of auctions with profit-motivated human subjects, often executed on a network of computers to implement the market rules, enables us to characterize the market outcomes of a variety of existing and new market designs.

Beyond the ability to compute what would otherwise be difficult or impossible, computer simulations can help us discover knowledge in a more important way.

> Artificial systems and adaptive systems have properties that make them particularly susceptible to simulation via simplified models. … Resemblance in behavior of systems without identity of the inner systems is particularly feasible if the aspects in which we are interested arise out of the *organization* of parts; independently of all but a few properties of the individual components. Thus for many purposes we may be interested in only such characteristics of a material as its tensile and compressive strength. We may be profoundly unconcerned about its chemical properties, or even whether it is wood or iron (Simon 1996, p. 16-17).

Computer simulations have served this role in helping us to analyze market artifacts and to discover and understand how, at the interface of their inner and outer environments, the elemental forces of want and scarcity interact through the laws of statistics. Simulations also reveal that a key property of fundamental concern in economics arises from the *organization* of its inner environment, largely independent of the decision-making behavior of individuals who constitute their outer environment. Let us turn to this discovery.

An Exploration with Zero Intelligence

Many investigative reports and the press blamed the stock market crash of October 1987 on program trading. Skeptical of such claims, I designed and taught a course on program trading at Carnegie Mellon University, hoping to learn in the process about the inner workings of double auction markets and the structure of trading strategies used in them.[1] Dhananjay Gode and I wrote

[1] In a double auction a buyer can submit a price at which he is willing to buy (make a bid), and a seller can submit a price at which he is willing to sell (make an ask). If another buyers bids a higher price, it becomes the market bid; if another seller asks a lower price, it becomes the market ask. A buyer is free to accept the market ask; a seller is free to accept a market bid and

double auction software (Market 2001, see Gode and Sunder 1994) for human as well as robot traders. Each student in the class could trade from the keyboard or could let a proxy trading strategy he wrote in the form of computer code replace him. Figure 2 (from Gode and Sunder 1994) shows the price paths generated in three different trading sessions with identical market demand and supply conditions. The feasible range for prices was 0-200, and the market demand and supply functions intersected in the price range 82-86.

The top panel of Figure 2 shows the time series of prices in a market where students traded among themselves under a promise that a part of their course grade would be proportional to the number of points earned by each trader. These data simply replicate the results of many classroom auctions with profit-motivated students conducted over the past half a century (e.g., Chamberlin 1948, Smith 1982, and Plott 1982). After some initial variability, double auction prices and allocations in classical market environments settle down in the neighborhood of the predictions of theory, even with a mere handful of traders. When subject rewards are not linked to the points earned in the auction, such markets still tend to settle down to the same predictions, albeit less reliably so (Jamal and Sunder 1991).

The second panel of Figure 2 shows the price series observed in a market in which each human trader had been replaced by an artificially intelligent robotic proxy in the form of a computer program written by the trader. In this market, the prices started higher, close to 100, in the middle of the price range 0-200, and

---

such acceptances consummate a binding transaction. The auction continues for a specified period of time.

settled down to a level slightly below the equilibrium range of 82-86. A significant amount of excess volatility persisted even after several periods. These programs seem to "learn" more slowly than their human progenitors; even after several periods, they make many more bids and offers per transaction than human traders do.

Examination of the student-written codes reveals a large variation across the artificially intelligent trading strategies. We cannot be sure what decision rules the human traders who wrote these codes used for trading with their fingers. From the individual bids and offers we can infer that the trading rules embedded in the computer codes were quite distinct. Business students found it difficult, both conceptually and technically, to express their intended trading strategies in the form of state-contingent and dynamically learning computers code. They pressed us for our own trading strategy so they could trade against it—and beat it.

Through several weeks of this program trading course, we had used the allocative efficiency of the markets (total profits earned by all traders as a fraction of the maximum total profits that could have been earned) as an index of the overall quality of students' coded trading strategies as they evolved after competing in successive class sessions. Until this point we had believed that as students learned better to formalize and translate their thinking into computer code, both price path and allocative efficiency should converge to their equilibrium values. After all, when the uninitiated students trade with their fingers, prices and efficiency of the markets come close to the equilibrium values within 10-15 minutes of trading. Within five weeks of the course, the allocative efficiency of markets with artificially intelligent traders crept up slowly from around 60 percent to 90 percent. Students' coded strategies were getting smarter, making fewer errors (their program taking an action the author had not intended under the circumstances). We thought it was just a matter of time before the codes would become as smart or smarter than their authors, and markets populated by them would achieve 100 percent allocative efficiency.

Meeting the students' challenge to the instructor for a coded strategy presented two problems. Gode and I knew little about what is a good trading strategy in a double auction (Wilson 1987); such learning itself had motivated the design and offering of this course. We wondered if our inability to write a winning strategy might raise questions about our suitability for teaching the course. We were not sure if beating the instructor's strategy would energize or demoralize the class; we knew it would demoralize us.

Toward the end of the term, we finally wrote a trading strategy. The bottom panel of Figure 2 shows the results of a market consisting entirely of clones of this program. The market demand and supply remain unchanged.

This market exhibits more variability in prices than the previous two. The strategy consists of one line of computer code: if you are a seller with a cost of, say, 40, pick a uniformly distributed random number between 40 and 200 and submit it as an "ask"; if you are a buyer with a value of, say 135, pick a uniformly distributed random number between 0 and 135 and submit it as a "bid." This strategy makes sure these traders—later labeled "zero-intelligence" or ZI traders—do not trade at a loss but keep spewing new proposals. No maximization, no memory, no learning, no natural selection, and no arbitrage. Yet, prices in this market also converge to a level near the equilibrium prediction of the neoclassical model.

Our motivation for the ZI strategy was part jest: it was sure to lose to the student strategies, but we could still save face with such an obviously simple and silly strategy. But it also arose from partially formed ideas after hours of watching the bids, asks and transaction prices of double auction trading on dynamically moving charts of Market 2001 computer screens. Human traders learn quickly enough for their markets to achieve almost 100 percent efficiency. Markets with artificially intelligent traders seem to fail only because these traders get stuck doing nothing in a contingency their authors had not anticipated. What might happen if the traders keep trying to trade without losing money? We did not know. Human experiments could not answer the question. Computers, with their ability to model the micro-level behavior of traders in any manner we want, helped us find out. And the answer surprised us.

Introduction to the elements of economics derives competitive equilibrium as an outcome of the individual striving to maximize their personal gain. From Adam Smith to the modern mathematical derivation of the first fundamental theorem of economics, this maximization is etched into our economics consciousness. In laboratory experiments with human traders that we, following others, have conducted, rewarding subjects on the basis of their performance to encourage them to maximize their rewards is an important part of the method.

Yet we found that prices in this market converged without any attempt by the traders to maximize.

Examination of the allocative efficiency of the markets held an even greater surprise for us: efficiency of the third market with zero-intelligence (ZI) traders is almost the same (about 99 percent) as the efficiency of the first market with profit-motivated human traders. We may not fully understand the decision rules of the human traders, but there is no mystery about the behavior of the ZI traders. We know for sure that they do not maximize; they are programmed merely to pick prices randomly with an opportunity set defined by a no-loss constraint. This is analogous to Becker's (1962) consumers whose random choices from their opportunity set in commodity space generate a downward-sloping demand function. The ZI traders, who bear little resemblance to their human counterparts in their motivation, cognitive equipment, or decision rules yield market outcomes that are virtually identical in allocative efficiency—the critical performance feature of the market artifact.

Simon had developed and validated the bounded rationality theory of individual decision-making decades earlier. These results suggest that the achievement of high levels of efficiency under classical conditions may place minimal demands on individual rationality—no maximization and not even bounded rationality is necessary. If individuals simply refrain from throwing their money away by making "obviously stupid" trades given their local information, allocative efficiency approaches its maximum.[2] After a decade of mathematical modeling, reprogramming of robots, analyses of data, and more simulations, what are the findings of this work.

Some features of market outcomes are largely robust to variations in the decision-making behavior of agents who participate in them. Allocative efficiency, a key measure of market outcomes, is one such feature. Adam Smith's conclusion that the allocative efficiency arises from individual pursuit of self-interest may be more general than it appears. Efficiency is achievable in double auction markets even if agents act randomly within their budget constraints. Random choice within one's opportunity set is, at best, only a weak form of "pursuit of self-interest" (Gode and Sunder 1993a).

The use of the maximization assumption to derive market equilibria in economics and the findings from cognitive psychology that individuals cannot and often do not know how to maximize need not be seen to be mutually inconsistent. Market institutions[3] may be the society's way of dealing with the human cognitive limitations. In classical environments, markets can approach the aggregate maximum even if the individuals do not know how to.

Efficiency of markets is primarily a function of their rules. Most of the efficiency arises from two basic rules: traders abiding with their proposals and priority for proposals disadvantageous to their originators (i.e., high bids and low asks). Contrary to the teachings of standard textbooks, the shapes of market

[2] John von Neumann (1956) points out that the link between the details of the components of a system to the performance of the system can be quite weak; only a few aspects of a component may be functionally relevant to the system.
[3] Also see North (1990).

demand and supply in extra-marginal region influence allocative efficiency (Gode and Sunder 1993b and 1997).

As the market demand and supply conditions change, the expected loss of efficiency has an upper bound. This bound is generated by a trade-off between the magnitude and the probability of efficiency loss associated with the displacement of intra- by extra-marginal traders. This market-level trade-off is independent of the individual trade-off between a proposal's profit and its probability of being accepted.

Double auctions are more efficient than one-sided auctions such as sealed-bid auctions because the former require more conditions to be fulfilled for an inefficient trade to occur. On one hand, auctions that batch or accumulate bids and asks before picking the highest bid and the lowest ask are more efficient than auctions where a transaction occurs as soon as a bid exceeds or equals an ask. Such auctions have lower probability of allowing the extra-marginal traders to displace the intra-marginal traders; other things being the same, call markets are favored over continuous auctions. On the other hand, efficiency is higher if traders can observe market data (e.g., call auctions in which the bids and asks are made public in real time, as compared to call markets in which they are not made public). Public bids and offers allow the intra-marginal traders to promptly outbid or undercut the extra-marginal traders, again reducing the probability of efficiency reducing displacement of intra-marginal traders.[4] In asset markets where the value itself is discovered through the market process, continuous

[4] Also see Cason and Friedman (1998).

markets have the advantage of faster price discovery; and the disadvantage of lower allocative efficiency.

Single-market findings about double auctions generalize to a set of multiple interlinked markets. If inventories are maintained between the markets, the effect of market discipline weakens and efficiency drops (Bosch and Sunder 2000). The partial equilibrium result on achievement of Pareto efficient outcomes is replicated with ZI traders in simple general equilibrium setting of Edgeworth's Box for two commodities (Gode, Spear and Sunder 2001).

Double auction asset markets with state uncertainty and imperfect information converge to the equilibrium derived by assuming that the traders are profit-maximizing Bayesians, irrespective of whether the traders are actually (1) Bayesians, (2) empirical Bayesians, or (3) biased heuristic traders who use adaptive heuristics well known to be biased (Jamal and Sunder 1996, 2001).

Walrasian tatonnement is a valuable static model that captures the asymptotic behavior of markets, but it does not organize the data from the process of arriving at equilibrium well. The ZI model is a simple model that does a reasonable job of capturing the dynamics of markets, and organizing the data from the early part of trading well. The two models, in combination, may do a better job than either can do alone in helping us understand the markets.

Conclusions

Markets are powerful social institutions. They probably evolved in human societies because their efficiency had survival value. We can usefully distinguish between the inner and outer environments of an artifact. The former are designed to obtain a degree of insulation across variations in the latter, so the artifact can serve the function for which it is created or used. The inner environment of markets is defined by their rules; their outer environment includes the behavior of agents.

A claim that the predictions of the first fundamental theorem in economics are approachable in classical environments without actual or attempted maximization by participants might have been met with skepticism until recently. Thanks to a largely serendipitous discovery using computer simulations of markets, we can claim that weak forms of individual rationality, far short of maximization, when combined with appropriate market institutions, can be sufficient for the market outcomes to approach the predictions of the first fundamental theorem. These individual rationality conditions (labeled zero-intelligence) are almost indistinguishable from the budget or settlement constraints imposed on traders by the market institutions themselves. They are even weaker than Simon's concept of bounded rationality.

ZI traders are only an important first step toward using computer simulations with artificially intelligent traders to explore the structural properties of markets. Such simulations—the "wind tunnels" of economics—have already given us interesting discoveries. For example, we now know that the market level trade-off between the level and the probability of execution of an ask would exist even if no trader included such a trade-off in his strategy. Much more remains to be done.

As social artifacts, markets are the arena for the interplay of demand and supply. Functionality of markets can be assessed by their robustness to certain environmental variations and responsiveness to others. We prefer markets to be robust to variations in individual cognitive capabilities and responsive to their wants and resources. If creation without a creator and designs without a designer are possible, we need not be surprised that markets can exhibit elements of rationality absent in economic agents.

References

Becker, Gary S. 1962. "Irrational Behavior and Economic Theory," Journal of Political Economy 70: 1-13.

Bosch-Domenech, Antoni, and Shyam Sunder. 2000. "Tracking the Invisible Hand: Convergence of Double Auctions to Competitive Equilibrium," Computational Economics 16:3 (December): 257-284.

Cason, T. and Daniel Friedman. 1998. "Price Formation and Exchange in Thin Markets: A Laboratory Comparision of Institutions," in P. Howitt, E. de Antoni and A. Leijonhufrud, eds., Money, Markets and Method: Essays in Honour of Robert W. Clower. Elgar.

Friedman, D. and J. Rust. 1993. The Double Auction Market, Santa Fe Institute Studies in the Sciences of Complexity, Proc. Volume XIV, New York: Addison-Wesley.

Friedman, Milton. 1953. "The Methodology of Positive Economics," in Essays in Positive Economics. Chicago: University of Chicago Press.

Gode, D. K. and Shyam Sunder. 1993a. "Allocative Efficiency of Markets with Zero Intelligence Traders: Market as a Partial Substitute for Individual Rationality," The Journal of Political Economy, Vol. 101, No. 1 (February): 119-137.

Gode, D. K. and Shyam Sunder. 1993b. "Lower Bounds for Efficiency of Surplus Extraction in Double Auctions," in D. Friedman and J. Rust, editors, The Double Auction Market: Institutions, Theories and Laboratory Evidence Santa Fe Institute Studies in the Sciences of Complexity. New York, NY: Addison-Wesley, 199-219.

Gode, D. K. and Shyam Sunder. 1994. "Human and Artificially Intelligent Traders in a Double Auction Market: Experimental Evidence," in Kathleen Carley and Michael Prietula, editors, Computational Organization Theory, Hillsdale, NJ: Lawrence Erlbaum Associates: 241-262.

Gode, D. K. and Shyam Sunder, "What Makes Markets Allocationally Efficient?," Quarterly Journal of Economics Vol. CXII No. 2 (May 1997), 603-630.

Gode, D. K. and Shyam Sunder. 2000. "Double Auction Dynamics: Structural Effects of Non-binding Price Controls," Yale University Working Paper, Revised March 2000.

Gode, D. K., Stephen Spear and Shyam Sunder. 2001. "Convergence of Double Auctions to Pareto Optimal Allocations in Edgeworth Box," Yale University Working Paper, Revised November 2001.

Hayek, Friederich A. 1945. The Uses of Knowledge in Society, The American Economic Review, 35 (September): 519-30.

Jamal, Karim and Shyam Sunder. 1991. "Money vs. Gaming: Effects of Salient Monetary Payments in Double Oral Auctions," Organizational Behavior and Human Decision Processes 49 (June): 151-166.

Jamal, Karim and Shyam Sunder. 1996. "Bayesian Equilibrium in Double Auctions Populated by Biased Heuristic Traders," Journal of Economic Behavior and Organization Vol. 31 No. 2 (November): 273-291.

Jamal, Karim, and Shyam Sunder. 2001. "Why do Biased Heuristics Approximate Bayes Rule in Double Auctions?" Journal of Economic Behavior and Organization, 46:4: 431-435.
Knight, F. H. 1947. “The Planful Act,” part VI, in *Freedom and Reform*, Harper & Brothers, New York.
Larkey, Patrick D. 2002. “Ask A Simple Question: A Retrospective on Herbert Alexander Simon,” Working Paper, Carnegie Mellon University.
Leijonhufvud, A. 1993. “Towards a Not-Too-Rational Macroeconomics”, Southern Economic Journal 60 (1): 1-13.
North, Douglass C. 1990. Institutions, Institutional Change, and Economic Performance. New York: Cambridge University Press.
Plott, C. R. 1982. “Industrial Organization Theory and Experimental Economics,” Journal of Economic Literature 20 (December): 1485-1527.
Plott, C. and V. Smith. 1978. “An Experimental Examination of Two Institutions”, *Review of Economic Studies*, 45: 133-153.
Rothschild, E. 1994. “Adam Smith and the Invisible Hand,” American Economic Review, 84, no. 2: 319-322.
Simon, Herbert A. 1947. Administrative Behavior. New York: The Free Press.
Simon, Herbert A. 1978. “Rational Decision Making in Business Organizations,” American Economic Review, 69: 493-513.
Simon, Herbert A. 1996. The Sciences of the Artificial. Third Edition. Cambridge, Mass.: MIT Press.
Smith, Vernon L. 1982. “Macroeconomic Systems as an Experimental Science,” American Economic Review 72 (December): 923-55.
Smith, Vernon L. 1991. “Rational Choice: The Contrast between Economics and Psychology,” Journal of Political Economy, 99, no. 4: 877-897.
von Neumann, John. 1956. “Probabilistic Logics and the Synthesis of Reliable Organisms from Unreliable Components,” in C. E. Shannon and J. McCarthy, eds. Automata Studies. Princeton: Princeton University Press.
Wilson, R. 1987. “On Equilibria of Bid-Ask Markets, in Arrow and the Ascent of Modern Economic Theory, G. Feiwell ed.: 375-414. MacMillan, New York.

Figure 1: A Social System as an Artifact with Inner and Outer Environments

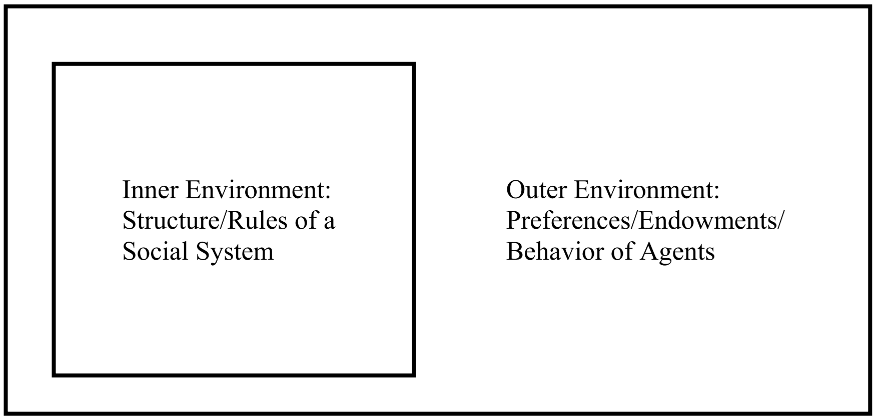

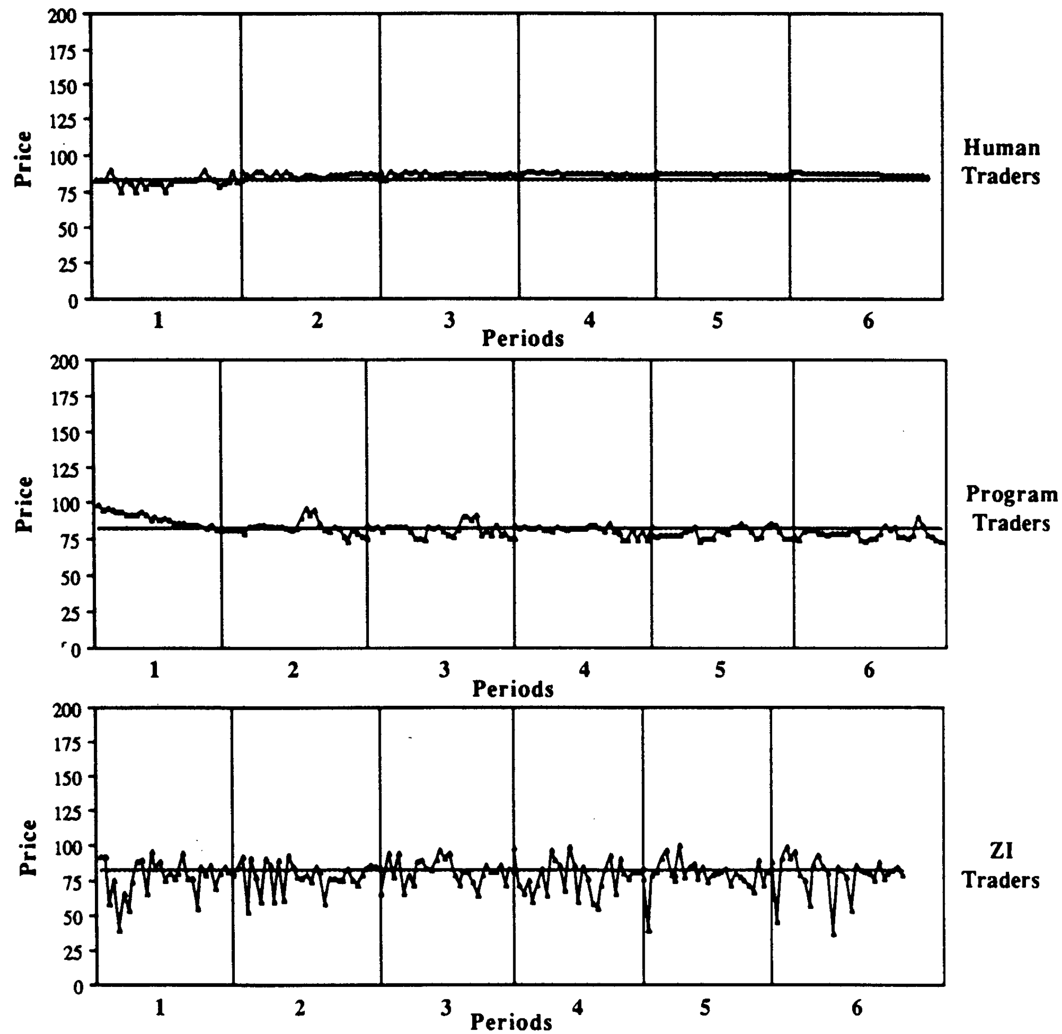

Figure 2: Price Series from Three Market Simulations
Source: Gode and Sunder (1994, p. 247).