Shop 'Til You Drop I: Market Trading Interactions as Adaptive Behavior

Dave Cliff*, Janet Bruten
Living the Vision
HP Laboratories Bristol
HPL-98-58
March, 1998

E-mail: davec@ai.mit.edu
jlb@hplb.hpl.hp.com

market, economics, trading, agent, adaptive behaviour

We argue that human economic interactions, particularly bargaining and trading in market environments, can be considered as adaptive behaviors. Moreover, the tools and techniques of adaptive behavior research could be profitably employed to build predictive models of existing or planned market systems. In addition to applications in economic modelling, "trading animats" could find use in market-based resource-allocation and control, and in internet-based commerce. Despite these potential applications, we note that there is a near-total absence of papers in the adaptive behaviour literature (and also in the artificial life literature) that deal with autonomous agents capable of exhibiting trading behaviors. After a brief overview of core concepts in microeconomics, we summarize work in experimental economics where human trading behavior is studied under laboratory conditions. We propose that such experiments could and should be used as 'benchmarks' for evaluating and comparing different architectures and strategies for trading animats. This paper is, essentially, a position paper: a manifesto calling the attention of the adaptive behavior research community to an entire field of problems that it has apparently so far ignored. In a companion paper [18], we present empirical results from simulations that invite a Braitenberg-style eliminative materialism perspective on the dynamics of experimental retail markets.

Internal Accession Date Only

*Artifical Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts

Shop 'til you drop I: Market Trading Interactions as Adaptive Behavior

Dave Cliff
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
CAMBRIDGE MA 02139, U.S.A.
davec@ai.mit.edu

Janet Bruten
Hewlett-Packard Laboratories Bristol
Filton Road
BRISTOL BS12 6QZ, U.K.
jlb@hplb.hpl.hp.com

Abstract

We argue that human economic interactions, particularly bargaining and trading in market environments, can be considered as adaptive behaviors. Moreover, the tools and techniques of adaptive behavior research could be profitably employed to build predictive models of existing or planned market systems. In addition to applications in economic modeling, "trading animats" could find use in market-based resource-allocation and control, and in internet-based commerce. Despite these potential applications, we note that there is a near-total absence of papers in the adaptive behavior literature (and also in the artificial life literature) that deal with autonomous agents capable of exhibiting trading behaviors. After a brief overview of core concepts in microeconomics, we summarize work in experimental economics where human trading behavior is studied under laboratory conditions. We propose that such experiments could and should be used as 'benchmarks' for evaluating and comparing different architectures and strategies for trading animats. This paper is, essentially, a position paper: a manifesto calling the attention of the adaptive behavior research community to an entire field of problems that it has apparently so far ignored. In a companion paper [18], we present empirical results from simulations that invite a Braitenberg-style eliminative materialism perspective on the dynamics of experimental retail markets.

1 Introduction

The majority of research in adaptive behavior has concentrated on developing artificial autonomous agents (i.e., animats) that exhibit behaviors common to many species of animals. Typical examples are spatiotemporal behavior patterns such as obstacle-avoidance, wall-following, and navigation to goal locations. Very few papers in the adaptive behavior literature have examined behaviors that are exhibited exclusively by humans.

In this paper we argue that human economic interactions, particularly bargaining and trading in market environments, can be considered as adaptive behavior despite being uniquely human. One strength of adaptive behavior research is its equal emphasis on synthesis and analysis, both in explaining behaviors of animals and in designing animal-like artefacts. We argue that this approach can be applied to human economic activity, and in a companion SAB paper [18] we describe experiments where simple adaptive animats interact within environments based on experimental retail markets used to evaluate human trading behavior. We show that the collective behavior of the trading animats is similar to that of the groups of humans, and that explanations of the animat markets could have significant impact on the way in which comparable human activity is explained.

If groups of simple artificial agents interact to exhibit market-level phenomena that are similar to those of human markets, explanations of how the phenomena arise in the artificial system may be viewed as candidate explanations for the same phenomena in human markets. Thus, adaptive behavior techniques can be used to build explanatory models of existing or planned market systems. It should also be possible to use such models for predictive purposes: at the policy level, the effects of possible changes in the organization of the market could be explored in simulation rather than by trial-and-(expensive)-error in the real world; at a more avaricious level, predictions of the future behavior of various markets offers manifest opportunities for generating income, so long as the predictions are accurate.

We doubt that current agent technology and human performance data could be combined to create genuinely novel explanatory or predictive models, but we see this as a worthwhile and challenging aim for future work. However, in addition to economic modeling (either theoretical or applied), there are two other significant uses in which artificial agents with human-like bargaining or trading behaviors could be employed: market-based control; and internet-based commerce.

In market-based control (MBC), economics is used as a source of inspiration, metaphors, and terminology for developing solutions to problems in distributed resource-allocation and control [36, 14]. In brief, the aim of MBC systems is that groups of software 'agents' or 'traders' interact within a market-like framework. The inspiration comes from free-market economics in the form of a division of the scarce resources into units of 'commodity' and the provision of a 'currency' that allows the agents to buy and sell the commodity. Some agents act as 'producers' or 'sellers' of the commodity (e.g., an agent may be assigned to a node or link in a telecommunications network, charging for use of that resource), while others act as 'consumers' or 'buyers' (e.g., an agent may be assigned to a data-packet on a network, spending currency in order to route the packet from its source to its destination). Crucially, as with real free markets, the MBC system should do this in a distributed fashion: there should be no central or global control process; rather, the allocation of resources 'emerges' from the local interactions of the buyers and sellers. With the growing interdependence and integration of networked computer systems and telecommunications systems, possibly involving multiple owners or vendors, MBC systems are increasingly being viewed as a highly viable means of decentralizing control: they offer potential for robustness and readily accommodate users with varying quality-of-service needs. Yet, to the best of our knowledge, no current MBC systems are fully decentralized and automated. In the applications published in the literature, there is a reliance either on centralized 'auctioneer' processes or on human intervention. Thus, there is a need for 'bargaining' mechanisms that allow software agents to agree on prices for transactions in decentralized MBC systems. For a more detailed critique of current MBC work, and further explanation of how our trading animats address the needs of MBC applications, see [15, 19].

The recent explosive increase of activity on the internet and world-wide-web, and in particular the announcement of secure methods for electronic transfer of funds, offers the potential for exploiting internet-based commerce. One way in which goods and services could be advertized and sold on the internet is the use of software agents that autonomously seek out and purchase items on behalf of a human user, entering into bargaining (or "haggling") interactions with agents representing the sellers of the desired items. The user would, presumably, indicate preferences such as price-range, desirable features, et cetera, and the buyer-agent would then traverse the web, seeking information from vendors' websites and interacting with each vendor's seller-agent, trying to get a good price. While the application of such techniques to the purchase of domestic goods such as music systems or kitchenware may be somewhat fanciful, a more realistic possibility might be the use of autonomous bargaining agents to automate the trading of commodities in firm-to-firm markets where the overall cost or volume of the transactions is not so great that it provides an incentive for the businesses concerned to engage in special relationships, such as approved quality control or bulk discounts. When the costs and volumes are sufficiently low, the identity of a seller may be irrelevant to the buyer, and the buyer is unlikely to be seeking a long-term supply-chain management relationship. If the trading is not part of a company's core business, but part of the fundamental "bread and butter" purchasing all large companies need to do, the identity of the seller is also likely to be less important. Markets exist for trading overstocked components or products in a number of industries, and some of these are moving to internet-based anonymous electronic trading. Typically in such markets the sellers will get a very poor price if they sell to a broker, but equally don't want to spend time and money trying to find a buyer. The buyers in these markets are often interested in small quantities, with very unpredictable demand, and thus are unlikely to be able to benefit from bulk discounts and special relationships with large producers. Such electronic markets offer may become significant application areas for autonomous trading agents.

Section 2 explains our arguments in more detail. We start, in Section 2.1, with a brief review of some fundamental issues in microeconomics, the branch of economics that deals with market behaviors. Section 2.2 then summarizes seminal work in *experimental economics*, where the trading behavior of small groups of humans in particular markets is studied under experimental conditions. The clarity of the experimental economics results, where theoretically predictable market dynamics emerge as a result of the collective behavior of the (human) trading agents, acts as one motivation for the development of adaptive behavior models of market activity. This issue is explored in more depth in Section 2.3, where we present our arguments for treating trading and bargaining behaviors as adaptive behavior. In Section 3 we discuss the near-total lack of relevant work in the published literature on adaptive behavior and the related field of artificial life. Although this is a "position paper", presenting methodological arguments, we demonstrate empirical work in a companion SAB paper [18] and in other publications [15, 16, 17, 19, 20].

2 Economics and Adaptive Behavior

2.1 Microeconomics

In the economics literature, a distinction is made between *microeconomics*, the study of the structure and dynamics of particular markets, and *macroeconomics*, the study of the structure and dynamics of entire economies. A microeconomist might study the market for butter, possibly just in terms of how changes in supply and demand for butter affect its price, or possibly in relation to the market for margarine. Meanwhile, a macroeconomist might study the relation between inflation and unemployment, possibly in relation to a particular government's welfare, taxation, or money-supply policies, or perhaps as part of a cross-government international comparative study. But even complex global economies, affected by tangles of multinational corporations and national governments are, in principle at least, rooted in the fundamental microeconomic activity of supplying commodities for which there is some demand. Therefore, in principle at least, there is a smooth continuum from elementary microeconomic analyses up to complex macroeconomic theories and models.

2.1.1 Basics

Microeconomics is often characterized as being primarily concerned with the allocation of scarce resources: "Microeconomics is a set of theories with one aim: to help us gain an understanding of the process by which scarce resources are allocated among alternative uses, and of the role of prices and markets in this process." [32, p.1].

More precisely, a *market* can be defined as "...a set of arrangements by which buyers and sellers are in contact to exchange goods or services" [6, p.32]. The quantity of a commodity (good or service) that buyers are prepared to purchase at each possible price is referred to as *demand*, and the quantity of a commodity that sellers are prepared to sell at each possible price is referred to as *supply*. In general, the greater the price of a commodity, the fewer buyers will be inclined to make a purchase, and so the quantity demanded reduces: thus, if we plot price as a function of quantity, the *demand curve* slopes downward. In contrast, the greater the price of a commodity, the more sellers are inclined to sell, and so the quantity supplied increases: on a plot of price as a function of quantity, the *supply curve* slopes upwards. From these considerations, it is clear that at high prices the quantity supplied may exceed the quantity demanded (i.e., there is a surplus, or *excess supply*), and at low prices the reverse may be true (giving a shortage, or *excess demand*). But, at some intermediate price, the quantity demanded is equal to the quantity supplied: this is the *equilibrium price*, which 'clears the market': graphically, the equilibrium price (and quantity) can be determined by the intersection of the supply and demand curves, as illustrated in Fig. 1. Throughout the rest of this paper, the equilibrium price will be represented by P_0 and the equilibrium quantity by Q_0. Some authors refer to Q_0 as the "clearing quantity".

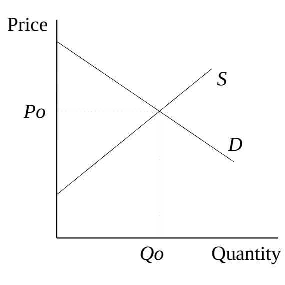

Figure 1: Simple illustration of supply and demand. The Supply Curve S slopes upwards and the Demand Curve D slopes downward. The two curves intersect at a point indicating the Equilibrium Price P_0 and the Equilibrium (Clearing) Quantity Q_0.

In theory at least, markets are *self-correcting*: if the supply and demand schedules remain fixed, the prices of transactions in the market (i.e., the *market price*) will tend toward P_0. At prices below P_0, there is excess demand: the suppliers can then choose only to sell to the highest-bidding buyers, and so the buyers have an incentive to bid higher prices, thereby raising the market price towards P_0. At prices above P_0, there is excess supply: buyers can then choose only to buy from the sellers with the cheapest offers, so sellers have an incentive to cut their offer prices, thus lowering the price towards P_0. When the market price is at P_0, neither buyers nor sellers have any incentive to change their prices, and so the system stabilizes. In this sense, the actions of a group of individuals in a market, each pursuing their own interests, can give rise to *price determination*, or *equilibration*, where the market price is P_0 and so the quantities demanded and supplied match at Q_0. Because the equilibrium is a result of price-competition between the agents in the market, it is often referred to as a *competitive equilibrium*. Such competitive market mechanisms, it is argued, can give efficient (or perhaps optimal) allocation of resources without centralized control or external intervention (e.g., by government regulation). A common ideal is *Pareto efficiency*: an allocation of resources is Pareto-efficient if no-one can be made better-off without someone else being made worse-off.

If there is no external intervention (e.g., price-caps imposed by some regulatory authority, such as a government), the system is said to be a *free market*. Free markets can give rise to 'emergent' collective behavior (convergence to competitive equilibrium) of the whole group which is in the best interests of that group (Pareto-efficient), despite the fact that each agent is operating only to satisfy self-interest. The group appears to be led by an 'invisible hand', in the metaphor introduced by Adam Smith in his 1776 book *The Wealth of Nations*. There are, however, conditions in which free-markets fail to achieve an efficient allocation: see [6, pp.264–265].

Conditions under which a market can attain equilibrium from a given set of initial conditions, and the nature of the approach to equilibrium, have been the subject of intense research in economics. Many models of equilibration assume *perfect competition*, where homogeneous indivisible units of a commodity are traded by large numbers of buyers and sellers, none of whom are sufficiently powerful to individually have any impact on market price, all of whom are aiming to maximize utility (i.e., sellers maximizing profits, buyers minimizing costs), and none of whom incur any costs in entering or leaving the market [5, p.325]. The nature or organization of some markets makes perfect competition unlikely or impossible, yet price determination can still occur. However, as the number of individuals in the market falls, the likelihood of collusion and the formation of rings or cartels increases: as the number of sellers is reduced, the market approaches an *oligopoly*, where the behavior of individual sellers in the market is highly dependent on the likely initiatives or responses of the other sellers. In such cases, the actions of individual buyers or sellers can have a significant impact on the market price, depending on the organization of the market.

2.1.2 Market Organization

Market organization concerns the regulations governing the information available to the agents in the market and the agents' *opportunity sets* (i.e., the possible actions they can perform). Together, these affect the method by which a market price is determined. In most markets, prices are determined via a particular style of *auction*. There are a large number of types of auction mechanism, and the literature on auctions is confused by the fact that different authors sometimes use the same name to refer to differing mechanisms: see [15] for a brief overview.

One particular style of auction that has received significant attention in the literature is the *continuous double auction*, or CDA. In a CDA, a group of buyers and a group of sellers simultaneously and asynchronously announce bids and offers: at any time, a seller is free to accept the bid of a buyer, and a buyer is free to accept the offer of a seller. One practical reason for the interest in CDAs is that they have for many years been the basis of trading in major international financial markets (such as the London and New York stock exchanges): originally as open-outcry oral auctions taking place on the trading floor of the exchange, and latterly as electronic auctions taking place in the cyberspace of city-wide networks of computerized dealing rooms.

While there are many other styles of auction,[1] much work has been concentrated on CDAs, motivated by the fact that they are often very fast and efficient, despite (or possibly because of) the volumes of information exchanged:

> "...markets organized under [CDA] trading rules appear to generate competitive outcomes more quickly and reliably than markets organized under any alternative set of trading rules. For this reason, [CDA] markets have been frequently investigated as a standard against which the performance of other institutions is evaluated." [23, p.126].

This interest has resulted in a research literature discussing CDAs (see, e.g., [29] and [23, pp.125–172] for reviews). The motivation for such work comes from a desire to better understand *how* the organization of the market, and the behavior of the traders in that market, affects the speed and efficiency of the market. This is summarized neatly by [49, p.156]:

> "Although the textbook 'supply equals demand' model may provide a good prediction of *closing* prices and quantities in [CDA] markets, it fails to explain the dynamics by which this happens. A more sophisticated theory is required to show how the trading process aggregates traders' dispersed information, driving the market towards [competitive equilibrium]. The ... problem was clearly stated by [(Hayek, *Amer. Econ. Rev.* 35(4):p.530, 1945)]: "The problem is in no way solved if we can show that all facts, *if* they were known to a single mind, would uniquely determine the solution; instead we must show how a solution is produced by the interactions of people each of whom possesses only partial knowledge. To assume all the knowledge to be given to a single mind in the same manner in which we assume it to be given to us as the explaining economists is to assume the problem away and to disregard everything that is important and significant in the real world."

Solving this problem, i.e. developing a theory of how the transaction prices of human traders converge on the theoretical equilibrium price, requires data gathered under controlled conditions. Such data has been generated by work in experimental economics, discussed next.

2.2 Experimental Economics

The use of 'laboratory methods' in economics, conducting controlled experiments to test theoretical hypotheses and predictions, has been of interest since at least the 1930's: for historical reviews, see [23, pp.5–9] and [48, pp.4–21]. In a typical simple experiment, a group of human subjects are each given the means to buy one unit of an arbitrary commodity; while another group are each given one unit of the commodity to sell. Each buyer will be a given a *limit price*, the maximum that buyer should pay for a unit of the commodity; and each seller will be given a minimum limit price below which the seller's unit should not be sold. Typically, different buyers will be given different limit prices, as will different sellers. The distribution of limit prices determines the supply and demand, and hence the values of P_0 and Q_0, for the experimental market. The subjects are then allowed to buy and sell within a particular market mechanism: in the early experiments, the markets were experimentally controlled open-outcry trading pits, but the vast majority of recent work has required the subjects to communicate via a computer network, which eases the control of free parameters and the recording of data. Theoretical hypotheses concerning the effects of a market's structure on its dynamics can be tested, and the implications explored by varying parameters such as supply and demand, the traders' opportunity sets, or the amount and quality of information available to each trader. Factors of interest may include the nature of the approach of observed transaction-prices towards P_0, the stability at equilibrium, the amount of potentially-available profit that is extracted from the market by the sellers, etc.

Davis and Holt [23] wrote a comprehensive text covering the major areas of experimental economics, while

[1] Examples include the *English* (ascending-bid), *Dutch* (descending-offer), *Vickrey* (second-price sealed-bid), *posted-offer*, and *call* auctions: see [15].

Kagel and Roth [39] edited a significant collection of critical surveys of the field. Methods and results from two key papers in the field are summarized below. Section 2.2.1 summarizes the first paper on experimental economics published by Vernon Smith, who helped establish the field and has since continued to be a prominent researcher: for a brief overview of his work, see [53], and for full details see his collection of papers spanning 30 years of research [52]. Section 2.2.2 then briefly summarizes work done by Gode & Sunder on 'zero-intelligence' traders that act randomly within CDA markets yet appear to give human-like market behavior.

2.2.1 Early Studies of Market Behavior

Smith [51] reported on experiments performed over a six-year period starting in 1955. All of the experimental regimes are similar to that described above: a cohort of human subjects are divided into a group of sellers and a group of buyers, and the two groups then trade within some specified market mechanism. Each trader's individual limit price is private (i.e., is not known by any other trader). Each buyer is encouraged to trade in the market by being instructed to consider the difference between the given limit price and the actual transaction price for the commodity as pure profit. Furthermore, buyers are told it is better to make no profit and own the commodity than to go without (i.e., they are encouraged to 'trade at the margin'). Similarly, each seller is told to treat the difference between the transaction price and the given limit price as profit, and to trade at the margin.

Each experiment is run as a sequence of distinct trading periods or 'days'. At the start of each day, all traders are allowed to quote a price: sellers quote *offers* (e.g., "sell at $2.50") and buyers quote *bids* (e.g., "buy at $1.20"). The quotes continue, typically with both groups of traders altering their quote-prices (increasing bids and decreasing offers) in an attempt at securing a transaction. At any time, a buyer is free to accept a seller's offer or a seller is free to accept a buyer's bid. When this happens, the buyer and seller are considered to have entered into a binding transaction: for both traders, the number of units they are entitled to trade in is reduced by one, and when a trader's entitlement reaches zero that trader drops out of the market for the remainder of that trading day. This process continues until the quotes of the traders no longer lead to contracts being made, or when some predetermined time-limit (typically five or ten minutes) is reached, at which point the day ends. If there are more days to run, the entitlements of all traders are then restored to their start-of-day values and the market reopens for another day's trading.

In a typical experiment, trading in the first day is characterized by early transactions taking place at prices that differ significantly from the P_0 value: as the day progresses, transaction prices approach P_0. On subsequent days, transaction prices are initially nearer P_0, and approach it faster. In most of the experiments described by [51], only the prices of agreed transactions were recorded.

To better characterize the convergence of transaction prices, Smith introduced a "coefficient of convergence", α, that is computed at the end of each day's trading in the market. This is defined as $\alpha = 100\sigma_0/P_0$, where, for a set of n transaction prices denoted by $P_{i:i\in\{1,\ldots,n\}}$, $\sigma_0 = (\frac{1}{n}\sum_{i=1}^{n}(P_i - P_0)^2)^{0.5}$. In most of Smith's experiments, α tends to decline from one trading day to the next. Smith also monitored the *allocative efficiency* of the market. This is defined as the total profit earned by all the traders divided by the maximum possible total profit that could have been earned by all the traders, expressed as a percentage. Typically, after one or two trading periods, human traders achieve allocative efficiency very close to 100%.

In the first eight experiments reported by Smith [51], each trader is allowed to buy or sell only one unit, although in later experiments this constraint was relaxed. Smith also experimented with changing the supply and demand during the experiment (i.e., after a few trading days, before the start of the next day's trading, a new set of limit prices were given to the subjects), and with having the buyers remain silent while only the sellers were allowed to quote offers: an experiment discussed in more detail in our companion paper [18]. Smith's results qualitatively indicated that the relationship of the supply and demand curves had an impact on the way in which transaction prices approached P_0: whether P_0 was approached from above or below, and whether the P_0 value was actually reached or the traders stabilized at a different value. Discussing these results in a subsequent publication, Smith states:

> "What have I shown? I have shown that with remarkably little learning, strict privacy, and a modest number [of traders], inexperienced traders converge rapidly to a competitive equilibrium under the [CDA] mechanism. The market works under much weaker conditions than had traditionally been thought to be necessary. You didn't have to have large numbers. Economic agents do not have to have perfect knowledge of supply and demand. You do not need price-taking behavior – everyone in the [CDA] is as much a price maker as a price taker." [52, p.157].

Smith's [51] results were some of the first to demonstrate that markets consisting of small numbers of traders could still exhibit equilibration to values predictable from classical microeconomic theory. To appreciate why this is significant, it is necessary to consider the underlying supply and demand curves in more detail.

Consider, for example, five sellers (denoted by the letters A to E) and five buyers (F to J) each with an entitlement to trade one unit. Possible supply and demand curves are illustrated in Fig. 2. From the figure, it is clear that P_0 is \$2.50, and Q_0 is three. Entitlements to buy or sell that determine the supply and demand curves to the left of the equilibrium quantity Q_0 are referred to as *intra-marginal* (or infra-marginal) units, while those determining the shape of the curves to the right of Q_0 are *extra-marginal*. Thus the entitlements of traders A, B, J, and I are intra-marginal, while D, E, F, and G are extra-marginal; C and H are simply *marginal*.

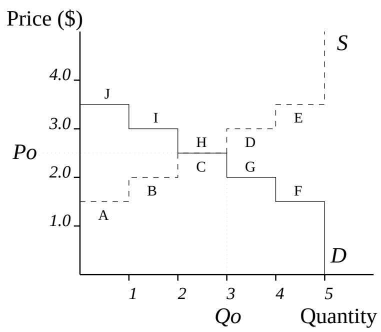

Figure 2: Supply and demand curves for ten traders. The supply curve is shown as a dashed line for extra clarity. There are five sellers (A to E), each with one unit to sell, and five buyers (F to J), each with the right to buy one unit. The step-changes in the quantities supplied and demanded are dependent on the limit prices of the individual traders, as indicated by the labels A to J. The intersection gives equilibrium values P_0=\$2.50 and Q_0=3.

As is evident in Fig. 2, the supply and demand curves for this experiment are stepped: this is because the commodity is dealt in indivisible discrete units, and there are only a small number of units available in the market. Thus, supply and demand in this simple market differs appreciably from the smoothly sloping curves of the idealized market illustrated in Fig. 1. The idealized market is based on conditions where the step-changes involved can be treated as infinitesimally small.

Furthermore, most classical theories of price determination and equilibration assume or require a large number of traders: if an individual trader drops out of the market, the supply and demand curves remain essentially the same. But in markets with few participants, this is not the case. Consider the simple market illustrated in Fig. 2: if the first trade in the market is between seller C and buyer I (at a price, say, of \$2.65: giving profits of \$0.15 for C and \$0.35 for I) then these two traders drop out of the market, but the resultant supply and demand curves, illustrated in Fig. 3, and the values of P_0 and Q_0 are significantly different.

Matters are further complicated when the difference between the *underlying* and *apparent* supply and demand curves is considered.[2] Each active trader in the

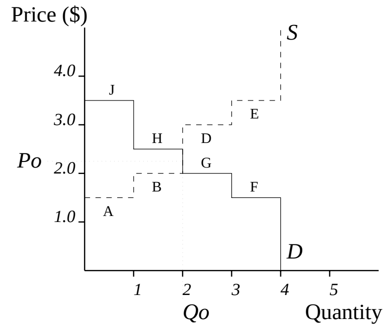

Figure 3: Supply and demand curves for eight traders. This is the market illustrated in Fig. 2, after traders C and I have agreed a transaction and left the market. The equilibrium price and quantity have altered. Now Q_0 is reduced to two, and P_0 is indicated at \$2.25, although technically it is no longer a scalar value: rather there is now a bounded range of possible equilibrium prices, from \$2.00 to \$2.50: what [23, p.131] refer to as a 'price tunnel'; any price within this range could be an equilibrium value.

market will be trying to make a profit, so buyers will be quoting prices lower than their individual limit prices, and sellers will be quoting prices higher than their limit prices. Because the limit price of each trader is private (i.e., known only to that trader), the prices quoted by the traders give only a weak indication of the *underlying* supply and demand determined by the limit prices: the *apparent* supply and demand, based on the actual prices quoted, may be significantly different. Thus, Figs 2 and 3 show underlying supply and demand before and after the trade between C and I, but an observer of (or participant in) the market does not have access to this information: these underlying schedules can only be guessed at by the traders, and the only information they have available is the quote-prices observed (i.e., "heard") in the market.

To illustrate the difference, assume that each trader aims for a particular 'profit margin', expressed as a percentage of the trader's limit price. Say that we take the market of Fig. 2 and assign each trader a randomly chosen profit margin in the range 0% to 50%. Then the market might have the apparent supply and demand curves illustrated in Fig. 4. As can be seen, the apparent supply and demand differ significantly from the underlying curves illustrated in Fig. 2. The values of P_0 and Q_0 are different, and the ranking of the traders' prices has altered.

Finally, it should be noted that the apparent supply and demand schedules are dynamic, and can alter rapidly. Because no trader has full knowledge of the underlying supply and demand, traders might base their profit levels on an initial guess that is then refined on the basis of the prices subsequently quoted by the competition (other members of the trader's group) and opposition (traders in the other group, or 'contraside' [8, p.226]), and on the basis of which quotes lead to transac-

[2] Smith [52, pp.809–810] refers to the underlying supply and demand as the *market* supply and demand, and to the apparent supply and demand as the *seller's offer array* and *buyer's bid array*: referred to collectively as the bid-and-offer arrays [23, p.300].

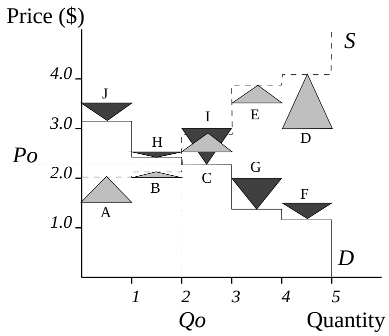

Figure 4: Bid-offer array for the ten-trader market illustrated in Fig. 2, given random profit margins between zero and fifty percent. Each buyer's limit and quote prices are illustrated as dark inverted triangles, while each seller's limit and quote prices are illustrated by light upright triangles: the base of each triangle indicates the trader's limit price, while the apex indicates the trader's quote-price. The array of bid-prices gives an apparent demand curve D, and the array of offer-prices gives an apparent supply curve S. The apparent supply and demand curves differ significantly from the underlying supply and demand shown in Fig. 2: see text for discussion.

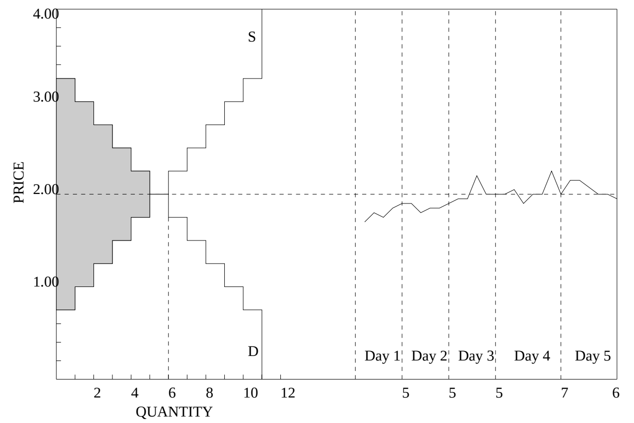

Figure 5: Redrawn from Smith's (1962) Chart 1: see text for discussion.

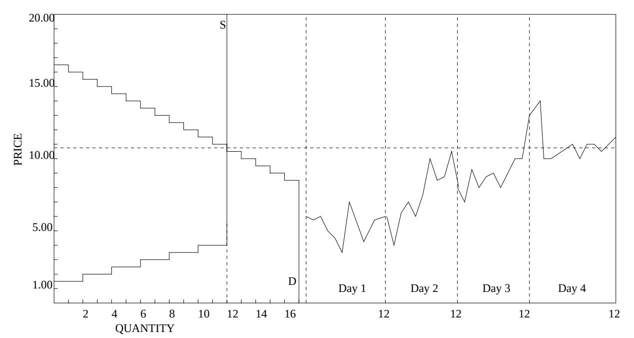

Figure 6: Adapted from Smith's (1962) Chart 6. See text for discussion.

tions and which are ignored. In a CDA, such information arrives in a continuous asynchronous stream. Moreover, the underlying supply and demand dynamically shift as traders enter and leave the market.

Despite this, the humans in Smith's experiments rapidly approach the competitive equilibrium predicted from theory. Figs 5 and 6 show some of Smith's [51] results. In both figures, the supply and demand curves are shown on the left, and the time series of transaction prices over a number of trading days is shown on the right. In Fig. 5, there are 11 buyers and 11 sellers, each with the right to buy or sell one unit: $P_0 = \$2.00$; $Q_0 = 6$. The shaded area in Fig. 5 indicates the available profit (or 'surplus' or 'rent') in the market: this is divided into two regions by the horizontal line at P_0, and Smith [51] hypothesized that the ratio of the areas of these two regions affected the nature of the approach of transaction prices to the theoretical equilibrium price. As is clear on the right of Fig. 5, transaction prices converge toward equilibrium over the five trading days, and the number of transactions per day varies from five to seven. In Fig. 6, there is excess demand (12 sellers but 17 buyers): transaction prices converge to equilibrium very slowly, and from below: when equilibrium is reached, there is evidence of some 'overshoot'.

Human beings are notoriously smart creatures: the question of just how much intelligence is required of an agent to achieve human-level performance is an intriguing one. This question was addressed by Gode & Sunder, whose work is summarized in the next section.

2.2.2 Zero Intelligence Traders

Gode & Sunder [30] describe a set of experiments similar in style to Smith's, but which use "zero intelligence" (ZI) programs that submit random bids and offers to replace human traders in electronic double-auction markets. They explore the performance of both 'unconstrained' and 'constrained' zero-intelligence traders, which they refer to using the abbreviations ZI-U and ZI-C, and compare the results of these traders to results from human traders operating in (almost) identical experimental conditions.

As with Smith's work, each ZI trader is given an entitlement to buy or sell a number of units, each with a particular limit price. The ZI traders generated quote prices at random: ZI-U traders generated quote-prices from a uniform distribution over the range of possible prices, regardless of their limit-prices. Thus, there was no constraint on the ZI-U traders to prevent them from entering into loss-making deals. In contrast, each ZI-C trader generated a quote-price for its current transaction using a uniform distribution that was constrained by that trader's limit-price for that transaction: this prevented the ZI-C traders from entering into loss-making deals.

Gode & Sunder [30] presented results from five experiments. In each experiment, a CDA market with specific supply and demand curves was run for a set number of trading days three times: once with ZI-U traders, once with ZI-C traders, and once with human traders. Significantly, and unexpectedly, the ZI-C traders appeared to give market performance (in terms of Smith's α measure, and in terms of allocative efficiency) that was much closer to the performance of human traders than it was to that of the ZI-U traders. In particular, the allocative efficiency of the ZI-C traders was virtually indistinguishable from that of the human traders. Thus, Gode & Sunder concluded that the market's attractive dynamics, i.e., low α and high allocative efficiency, is not much affected by the intelligence, rationality, learning, or memory capabilities of the traders. Rather, it is due to the structure of the market (i.e., traders interacting via a CDA and prevented from loss-making transactions). It appears that, so long as the market is structured correctly, traders with *no intelligence at all* can give human-like market dynamics.

Gode & Sunder's paper has had significant impact in the experimental economics literature;[3] it is mentioned in the current edition of Simon's classic book *The Sciences of the Artificial* [50, p.32]; and has even been discussed in Clark's latest book on the philosophy of cognitive science [13, pp.183–184].

However, we have recently presented analytic and empirical results which conclusively demonstrate that Gode & Sunder's results only occur in special circumstances, and that in general the ZI traders fail to give human-like market performance [15, 17, 20]. Thus, more than zero 'intelligence' is required of artificial trading agents to give human-like collective behavior and market dynamics. And because of this, methods developed in adaptive behavior research should be relevant.

2.3 Trading as Adaptive Behavior

Given that nonzero 'intelligence' is required of artificial traders to give human-like market dynamics, one possibility is to use traditional artificial intelligence (AI) techniques, based on deliberative reasoning with explicit representations of 'knowledge' or 'beliefs', as the foundation for building artificial trading agents. However, we believe that the assumptions underlying adaptive behavior research imply that it is a more attractive approach.

Although precise definitions are rarely articulated, much adaptive behavior research is consistent with the definition that a behavior is adaptive if, when an agent exhibits that behavior, it increases its chances of survival (or reduces its chances of ceasing to exist). This requires that the behavior is well-matched to the prevailing environmental conditions, and this carries an implicit assumption that the environment with which the agent interacts is nontrivial: i.e. that it is, to significant extents: complex, dynamic, unknown, unpredictable, and unforgiving of mistakes. Because of this, the agents studied in adaptive behavior are fundamentally *situated*, in the sense that their (inter)actions cannot be premeditated because it is impossible to predict with any accuracy all possible future conditions, events, or outcomes. Such situatedness requires that the agents are responsible for coordinating perception and action for extended periods of time without human intervention (i.e., that they are *autonomous*). Finally, adaptive behavior research typically places strong emphasis on the resource-limited nature of real-world interactions: agents in real environments do not have unlimited computational power or memory space or time in which to decide on appropriate actions. The stark contrast between these assumptions and those of the vast majority of work in AI is manifest.

Compare this explicit characterization of adaptive behavior with the situation faced by human traders in a real CDA market or in one of Smith's experiments: traders need to coordinate perception and action in the sense that they must accumulate and integrate data from diverse asynchronous sources, and act upon that data in good time (the value or integrity of much market data decays rapidly with time). In a competitive market, the environment is clearly dynamic and unforgiving. Relevant information (such as other traders' profit margins, or the information held by other traders) is rarely known or predictable, and it is unlikely that any trader will experience acts of kindness or selfless altruism. In real markets, traders that consistently fail to make profits will not last long. To summarise: if you're a trading agent, you can't be sure of anything except that the world is out to get you and that if you snooze you lose. For these reasons, trading agents need to be situated and autonomous. And, as adaptive behavior research is fundamentally concerned with the analysis or design of situated autonomous agents, there should be some use for adaptive behavior tools and techniques in studying or creating microeconomic systems. For instance, much work in adaptive behavior has explored the use of adaptation methods such as reinforcement learning (based on attempting to maximize "reward") and artificial evolution (based on attempting to maximize "fitness"). When trading agents interact within a market environment, measures of reward or fitness can be clearly and unambiguously equated with profit or utility.

Much work in microeconomics has been devoted to analytic studies (typically reliant on set-theory, differ-

[3] For example, the following texts approvingly cite Gode & Sunder's work: [8, pp.230–231], [10, pp.253, 258], [28, p.19], [29, p.*xxiii*], [40, pp.292, 294], [49, pp.160–161, 175], [23, p.132], [41, p.2], [33, p.1082], [54, p.310], [1, p.186], [9, p.674] [34, p.228], [35, p.370], [38, pp.570, 580], [42, p.226], [48, pp.52–55, 80–81], [57, p.475], [11, pp.1318, 1333], [21, p.678], [22, p.383], [37, p.276], [44, p.266], [56, p.2], [59, p.461], [12, p.320], [27, p.623], and [31, pp.604–605].

ential calculus, optimization methods such as dynamic programming, and game-theory) that are based on a number of simplifying assumptions necessary to maintain tractability. Simulations of the sort common in adaptive behavior research could be built: these would represnt models that are too complex for formal analytic treatment but sufficiently simple (relative to real systems) that they have useful predictive or explanatory power.

Given these arguments, it seems reasonable to expect that some prior work in the adaptive behavior literature will have explored issues in the collective behavior of agents that bargain and trade in market-based enviornments. Yet, surprisingly, the field seems to have almost totally ignored such issues.

3 (The Lack Of) Related Work

There are 256 papers published in the proceedings of the first four SAB conferences, and (at time of writing), a little over 60 papers published in the *Adaptive Behavior* journal. Given the arguments of the previous section, it would seem reasonable to expect that a fair proportion of papers in the adaptive behavior literature should deal with economic behaviors.

However, as far as we are aware, there is only *one* paper that deals explicitly with economic activity: this is Beltratti and Margarita's study of the evolution of trading strategies among heterogeneous artificial economic agents [7]. Beltratti and Margarita worked with artificial 'stock markets' populated by trading agents with neural network controllers that determined their trading strategies. In markets seeded with three initial types of strategy ('smart', 'dumb', and 'naive'), the dumb strategies rapidly disappeared, leaving 'smart' and 'naive' to interact. The agents could profit by making accurate predictions of the price of stocks: naive agents made predictions based only on the most recent market price, while the predictions of the smart agents required more sophisticated information that was provided at varying levels of cost. Beltratti and Margarita found that the market dynamics were dependent on the cost of information supplied to the smart agents. While this work is novel and interesting, it is of very little relevance to price-determination or bargaining behaviors, because whenever two agents entered into a transaction, the price was (arbitrarily) set at the average of the predictions of the two agents [7, p.495]. Thus, it appears that no work published in the adaptive behavior conference proceedings or journals addresses the issues we are discussing here.[4]

Adaptive behavior research is often characterized as related to work in *artificial life* (a-life). Many papers in the a-life literature deal with systems that exhibit complex coherent global behavior arising from the interaction of groups of components which are individually simple in comparison to their global behavior. Given this focus, the price-equilibration of groups of traders operating in market-based environments would seem to be a natural candidate for a-life research.

However, the international a-life journal and conference proceedings show a distinct lack of such research. While there is a small core of work on the *iterated prisoner's dilemma*, a classic game-theory problem in which the emergence of cooperative behavior among non-altruistic agents can be explored, and which is of direct relevance to oligopolistic markets (see, e.g., [4], [55], and [43]), we know of only two papers published in the a-life literature that explicitly study market trading strategies: [47] and [24]. Both of these papers report on the application of elementary evolutionary adaptation methods to optimize simple trading strategies for speculative markets, and both directly set all transaction-prices to be equal to P_0 via a centralized process that collects quote prices from all individuals, determines the supply and demand, and then calculates P_0: [47, p.191], [24, p.326]. For this reason, neither of these two papers are relevant to the study of equilibration or bargaining behaviors.

The apparent lack of work in a-life on agents with bargaining behaviors for market-based environments is confirmed in a recent review paper by Tesfatsion, [58]. The paper presents a summary overview of aspects of a-life especially relevant for the study of decentralized market economies. The two main areas of research activity discussed are: the combination of evolutionary game theory (e.g., iterated prisoner's dilemma) with preferential partner selection (i.e., the ability to choose or refuse particular opponents in the game); and an extension of this, where trade networks can form and evolve, using the prisoner's dilemma game to model risky trades between individuals. Thus, there is no emphasis on bargaining mechanisms in this work, and the indications from [58] are that little or no work in a-life is comparable to the work on trading agents discussed here.

Perhaps even more surprisingly, there appears to be very little relevant work in the economics literature. In a 1993 paper, Arthur [2] discussed the idea of designing artificial agents that behave like human economic agents, and explored the use of a simple classifier system as an adaptation mechanism for developing artificial agents that could be calibrated against human learning data from experiments exploring the psychology of economic behaviors. The use of artificial agents in economic modeling offers two benefits. First, the costs associated with using humans are eliminated. Second, the use of computer simulations forces a degree of mechanistic rigour, and thus helps clarify what information and/or cognitive mechanisms are necessary and sufficient to produce the behaviors of interest. In this sense, the algorithms or programs that specify the behavior of the artificial agents can be viewed as 'theories' concerning the gener-

[4] McFarland's work (e.g.,[45, 46]) uses microeconomics as a framework for comparing adaptive behavior in animals and robots, but does not explicitly address human economic behavior.

ation of comparable behaviors in real agents. However, Arthur's 1993 paper did not deal with price-bargaining in auctions.

Subsequently, Arthur collaborated with Holland, LeBaron, Palmer, & Tayler [3] to develop a system where autonomous software trading agents interact in an asset market, each using a classifier system to adapt its trading strategy to changes in the market. The agents adapted their 'expectations' of future price movements, and Arthur et al demonstrated that the market dynamics were affected by the rate at which alternative expectations were explored: at low rates of exploration, the market settled to a predictable equilibirum, but higher rates led to complex dynamics, including temporary bubbles and crashes.

We know of only two other papers comparable to the trading-agent work we advocate here: there are by Easley and Ledyard [25], and Rust et al [49]. We discuss these papers in detail in [15, 20], and also explain why other work, such as Epstein and Axtell's recent book [26] is not relevant. The significant issue here is that, to the best of our knowledge, there are very few papers in the economics literature that bear any worthwhile comparison to the work we propose here.

4 Conclusion

It is not clear to us why the published research in adaptive behavior, and also in artificial life, has largely ignored economic activity. We find it curious, given that there are strong reasons for considering economic activity as adaptive behavior. For surely, adaptive behavior research methods could (and should) be used to develop animats that engage in economic interactions. In addition to being used as potential scientific models of human activity, such animats could find use in several potential application areas including predictive simulations, market-based control, and internet-based commerce. The clear correlation between profitability (or utility-maximization) and 'fitness' or 'reward' should allow for many of the tools and techniques of adaptive behavior research to be readily transferred to autonomous agents that interact with market environments. Like many environments in adaptive behavior research, markets can be dynamic, uncertain, and unforgiving.

Given that the simplified and constrained market environments introduced by Smith have formed the basis of much work in experimental economics over the last 30 years, it seems wise to commence by attempting to develop "intelligent" trading agents with the bargaining capabilities necessary for small groups of traders to show price-equilibration in market scenarios similar to those studied with human subjects in experimental economics. By concentrating on a defined class of environments, the sharing of techniques and comparison of data should be made more easy, just as the development of the *Khepera* robot has provided a *de facto* standard platform for robotics-based adaptive behavior research.

The full details of equilbration in markets of human traders are yet to be explained. The development of trader animats capable of price-determination as a collective behavior in market-based environments could advance our understand of existing microeconomic systems and open up new application areas. Thus, we anticipate a melding of research in adaptive behavior and market microeconomics that we expect to be highly productive, and hopefully also highly profitable.

References

[1] M. J. Anderson & S. Sunder. Professional traders and intuitive bayesians. *Organizational Behavior and Human Decision Processes*, 64(2):185–202, 1995.

[2] W. B. Arthur. On designing economic agents that behave like human agents. *Evolutionary Economics*, 3:1–22, 1993.

[3] W. B. Arthur, J. H. Holland, B. LeBaron, R. Palmer, & P. Tayler. Asset pricing under endogenous expectations in an artificial stock market. In W. B. Arthur, S. Durlauf, & D. Lane, editors, *The Economy as an Evolving Complex System II*, pp.15–44. Addison Wesley, 1997.

[4] R. Axelrod. *The Evolution of Cooperation.* Basic Books, 1984.

[5] G. Bannock, R. E. Baxter, & E. Davis. *The Penguin Dictionary of Economics.* Penguin, fifth edition, 1992.

[6] D. Begg, S. Fischer, & R. Dornbusch. *Economics.* McGraw-Hill, fourth edition, 1994.

[7] Andrea Beltratti & Sergio Margarita. Evolution of trading strategies among heterogeneous artificial economic agents. In J.-A. Meyer, H. Roitblat, & S. Wilson, eds, *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior (SAB92)*, pp.494–501. MIT Press, 1992.

[8] T. Bollerslev and I. Domowitz. Some effects of restricting the electronic order book in an automated trade execution system. In [29, p.221–252], 1992.

[9] C. Camerer. Individual decision making. In J. H. Kagel & A. E. Roth, eds, *The Handbook of Experimental Economics*, pp.587–703. Princeton University Press, 1995.

[10] T. Cason & D. Friedman. An empirical analysis of price formation in double auction markets. In [29, pp.253–284], 1992.

[11] T. Cason & D. Friedman. Price formation in double auction markets. *Journal of Economic Dynamics and Control*, 20(8):1307–1337, 1996.

[12] T. Cason & D. Friedman. Price formation in single call markets. *Econometrica*, 65(2):311–345, 1997.

[13] A. J. Clark. *Being There: Putting Brain, Body, and World Together Again.* MIT Press, 1997.

[14] S. H. Clearwater, ed. *Market-Based Control: A Paradigm for Distributed Resource Allocation.* World Scientific, 1996.

[15] D. Cliff. Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical Report HP–97–91, Hewlett Packard Laboratories, Bristol, England, 1997.

[16] D. Cliff & J. Bruten. Less than human: Simple adaptive trading agents for CDA markets. Technical Report HP–97–155, Hewlett Packard Laboratories, Bristol, England, 1997. To be presented at CEFEES98, Cambridge UK, July 1998.

[17] D. Cliff & J. Bruten. More than zero intelligence needed for continuous double-auction markets. Technical Report HP–97–157, Hewlett Packard Laboratories, Bristol, England, 1997. To be presented at CEFEES98, Cambridge UK, July 1998.

[18] D. Cliff & J. Bruten. Shop 'til you drop II: Collective adaptive behavior of simple autonomous trading agents in simulated retail markets. Technical Report HP–97–1nn, Hewlett Packard Laboratories, Bristol, England, 1997.

[19] D. Cliff & J. Bruten. Simple bargaining agents for decentralized market-based control. Technical Report HP–97–1nn, Hewlett Packard Laboratories, Bristol, England, 1997.

[20] D. Cliff & J. Bruten. Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets. Technical Report HP–97–141, Hewlett Packard Laboratories, Bristol, England, 1997.

[21] J. Conlisk. Why bounded rationality? *Journal of Economic Literature*, 34(2):669–700, 1996.

[22] J. C. Cox & D. M. Grether. The preference-reversal phenomenon: Response-mode, markets, and incentives. *Economic Theory*, 7(3):381–405, 1996.

[23] D. D. Davis & C. A. Holt. *Experimental Economics*. Princeton University Press, 1993.

[24] M. de la Maza & D. Yuret. A futures market simulation with non-rational participants. In R. A. Brooks & P. Maes, eds, *Artificial Life IV*, pp.325–330. MIT Press, 1994.

[25] D. Easley & J. Ledyard. Theories of price formation and exchange in double oral auctions. In [29, pp.63–98], 1992.

[26] J. M. Epstein & R. Axtell. *Growing Artificial Societies: Social Science from the Bottom Up*. MIT Press, 1996.

[27] D. A. Evans. The role of markets in reducing expected utility violations. *Journal of Political Economy*, 105(3):622–636, 1997.

[28] D. Friedman. The double auction market institution: A survey. In [29, pp.3–26], 1992.

[29] D. Friedman & J. Rust, eds. *The Double Auction Market: Institutions, Theories, and Evidence*. Addison-Wesley, 1992.

[30] D. K. Gode & S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.

[31] D. K. Gode & S. Sunder. What makes markets allocatively efficient? *Quarterly Journal of Economics*, 112(2):603–630, 1997.

[32] H. Gravelle & R. Rees. *Microeconomics*. Longman, second edition, 1992.

[33] D. M. Grether. Individual behavior and market performance. *American Journal of Agricultural Economics*, 76(5):1079–1083, 1994.

[34] W. Hardle & A. Kirman. Nonclassical demand: a model-free examination of price-quantity relations in the Marseilles fish market. *Journal of Econometrics*, 67(1):227–257, 1995.

[35] C. A. Holt. Industrial organization: A survey of laboratory research. In J. H. Kagel & A. E. Roth, eds, *The Handbook of Experimental Economics*, pp.349–443. Princeton University Press, 1995.

[36] B. A. Huberman, ed. *The Ecology of Computation*. Elsevier/North-Holland, 1988.

[37] K. Jamal & S. Sunder. Bayesian equilibrium in double auctions populated by biased heuristic traders. *Journal of Economic Behavior and Organization*, 31(2):273–291, 1996.

[38] J. H. Kagel. Auctions: A survey of experimental research. In J. H. Kagel & A. E. Roth, eds, *The Handbook of Experimental Economics*, pp.501–585. Princeton University Press, 1995.

[39] J. H. Kagel & A. E. Roth, eds. *The Handbook of Experimental Economics*. Princeton University Press, 1995.

[40] J. H. Kagel & W. Vogt. Buyer's bid double auctions: Preliminary experimental results. In [29, pp.285–306], 1992.

[41] A. Leijonhufvud. Towards a not-too-rational macroeconomics. *Southern Economic Journal*, 60(1):1–13, 1993.

[42] G. Y. Luo. Evolution and market competition. *Journal of Economic Theory*, 67(1):223–250, 1995.

[43] R. M. May, S. Bohoeffer, & M. A. Nowak. Spatial games and the evolution of cooperation. In F. Morán, A. Moreno, J. J. Merelo, & P.Chacón, eds, *Proc. 3rd Euro. Conf. Artificial Life (ECAL95)*, pp.749–759, 1995. Springer-Verlag.

[44] R. P. McAfee & J. McMillan. Competition and game-theory. *Journal of Marketing Research*, 33(3):263–267, 1996.

[45] D. McFarland. What it means for robot behaviours to be adaptive. In J.-A. Meyer & S. W. Wilson, eds, *From Animals to Animats: Proc. 1st Int. Conf. Simulation of Adaptive Behavior (SAB90)*, pp.22–28, MIT Press, 1990.

[46] D. McFarland & T. Boesser. *Intelligent Behavior in Animals and Robots*. MIT Press, 1993.

[47] C. Nottola, F. Leroy, & F. Davalo. Dynamics of artificial markets: Speculative markets and emerging 'common sense' knowledge. In F. J. Varela & P. Bourgine, eds, *Proc. 1st Euro. Conf. Artificial Life (ECAL91)*, pp.185–194. MIT Press, 1991.

[48] A. E. Roth. Introduction to experimental economics. In J. H. Kagel & A. E. Roth, eds, *The Handbook of Experimental Economics*, pp.3–109. Princeton University Press, 1995.

[49] J. Rust, J. Miller, & R. Palmer. Behavior of trading automata in a computerized double auction market. In [29, pp.155-198], 1992.

[50] H. A. Simon. *Sciences of the Artificial*. MIT Press, Cambridge MA, third edition, 1996.

[51] V. L. Smith. An experimental study of competitive market behavior. *Journal of Political Economy*, 70:111–137, 1962.

[52] V. L. Smith. *Papers in Experimental Economics*. Cambridge University Press, 1992.

[53] V. L. Smith & A. Williams. Experimental market economics. *Scientific American*, 267(6):72–77, December 1992.

[54] D. O. Stahl & P. W. Wilson. Experimental evidence on players model of other players. *Journal of Economic Behavior and Organization*, 25(3):309–327, 1994.

[55] E. A. Stanley, D. Ashlock, & L. Tesfatsion. Iterated prisoner's dilemma with choice and refusal of partners. In C. G. Langton, ed, *Artificial Life III*, pp.131–176, Addison-Wesley, 1993.

[56] K. Steiglitz, M. Honig, & L. Cohen. A computational market model based on individual action. In [14, pp.1–27], 1996.

[57] S. Sunder. Experimental asset markets: A survey. In J. H. Kagel & A. E. Roth, eds, *The Handbook of Experimental Economics*, pp.445–500. Princeton University Press, 1995.

[58] L. Tesfatsion. How economists can get Alife. In W. B. Arthur, S. Durlauf, & D. Lane, eds, *The Economy as an Evolving Complex System II*, pp.533–564. Addison Wesley, 1997.

[59] M. V. Van Boening & N. T. Wilcox. Avoidable cost: Ride a double-auction roller-coaster. *American Economic Review*, 86(3):461–477, 1996.