Shop 'til you drop II: Collective Adaptive Behavior of Simple Autonomous Trading Agents in Simulated 'Retail' Markets

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

Phone +1 617 253 6625
Fax: +1 617 253 8080

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

Phone +44 117 922 8700
Fax: +44 117 922 9657

Abstract

In a companion paper [9] we argued that human economic interactions, particularly bargaining and trading in market environments, can be considered as adaptive behaviors, and that the tools and techniques of adaptive behavior research can be profitably employed in modeling naturally-occurring markets or constructing artificial market-based systems. If groups of simple artificial agents interact to exhibit market-level behaviors that are similar to those of human markets, explanations of how the behaviors arise in the artificial system may be viewed as candidate explanations for the same behaviors in human markets. In this paper, we illustrate these arguments by means of an example. We present results from experiments where an elementary machine learning technique endows simple autonomous software agents with the capability to adapt while interacting via price-bargaining in market environments. The environments are based on artificial retail markets used in experimental economics research. We demonstrate that groups of simple agents can exhibit human-like collective market behaviors. We note that, while it is often tempting to offer explanations of human market behavior in terms of the mental states of the agents in the market, our agents are sufficiently simple that mental states can have no useful role in explaining their activity. Thus, explanations of the human-like collective market behavior of our agents cannot be phrased in terms of mental states; thereby inviting comparisons with Braitenberg's influential "law of uphill analysis and downhill invention", with eliminative materialism in the philosophy of cognitive science, and with dynamical-systems-based analyses of adaptive behavior.

Topic-Areas: Collective and Social Behaviors; Software Agents; Applied Adaptive Behavior.

Shop 'til you drop II: Collective Adaptive Behavior of Simple Autonomous Trading Agents in Simulated 'Retail' Markets

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

In a companion paper [9] we argued that human economic interactions, particularly bargaining and trading in market environments, can be considered as adaptive behaviors, and that the tools and techniques of adaptive behavior research can be profitably employed in modeling naturally-occurring markets or constructing artificial market-based systems. If groups of simple artificial agents interact to exhibit market-level behaviors that are similar to those of human markets, explanations of how the behaviors arise in the artificial system may be viewed as candidate explanations for the same behaviors in human markets. In this paper, we illustrate these arguments by means of an example. We present results from experiments where an elementary machine learning technique endows simple autonomous software agents with the capability to adapt while interacting via price-bargaining in market environments. The environments are based on artificial retail markets used in experimental economics research. We demonstrate that groups of simple agents can exhibit human-like collective market behaviors. We note that, while it is often tempting to offer explanations of human market behavior in terms of the mental states of the agents in the market, our agents are sufficiently simple that mental states can have no useful role in explaining their activity. Thus, explanations of the human-like collective market behavior of our agents cannot be phrased in terms of mental states; thereby inviting comparisons with Braitenberg's influential "law of uphill analysis and downhill invention", with eliminative materialism in the philosophy of cognitive science, and with dynamical-systems-based analyses of adaptive behavior.

1 Introduction

In the companion paper [9] we argued that human trading interactions in market environments can be considered as instances of adaptive behavior. To illustrate this, we gave an overview of Smith's [21] seminal work in experimental economics, where human traders interact within a given market mechanism under 'laboratory' conditions. Smith's work was one of the first demonstrations that the transaction prices of small numbers of traders, interacting via a continuous double auction (CDA) market, could rapidly and reliably approach the theoretical equilibrium price, with no need for a centralized 'auctioneer'.

We noted that traders in such markets are autonomous and situated, and that, because adaptive behavior research is fundamentally concerned with autonomous situated agents – either real (animals) or artificial (animats) – the problem of creating artificial trading agents should no longer be ignored by adaptive behavior research. If successfully developed, 'trader animats' could be used both in the science of explaining human market activity and in the engineering of new microeconomic systems such as for internet-based commerce (e.g., [16, 18]) and market-based control (e.g., [5]). In all three cases, but especially in developing scientific models of human economic activity, significant amounts of further research are likely to be necessary before genuinely useful or productive systems can be created.

Although it may seem intuitively obvious that some form of 'intelligence' or adaptation is necessary in bargaining agents, Gode and Sunder [15] presented results that appear to indicate that their *zero-intelligence* (ZI) agents can exhibit human-like behavior in CDA markets. Gode and Sunder's ZI trading agents simply generated random prices for bids or offers, subject to the constraint that they could not enter into loss-making deals. However, we demonstrated [6, 8] that Gode and Sunder's result only holds in very specific circumstances and that, in general, some 'intelligence' in the form of adaptivity or sensitivity to previous and current events in the market is necessary. Hence, we give our trading agents adaptive capabilities by employing elementary machine-learning techniques. Because our agents are intended to have minimal intelligence, but not zero intelligence, they are referred to as "ZIP" traders: ZIP is an acronym for "zero-intelligence-plus".

In other publications [6, 7, 10, 11] we have shown that our ZIP traders do not suffer from the failings that afflict Gode and Sunder's ZI traders. Furthermore, we have noted that the collective behavior of groups of ZIP traders is *human-like*, by which we mean that ZIP traders in experimental CDA markets give scores on the standard metrics of market performance (such as the Smith's α measure of price-convergence [21], allocative efficiency, and profit dispersion [15]) that are very similar to those given by human traders in the same markets.

In this paper, we demonstrate that the market behavior of ZIP traders is human-like in another sense: ZIP traders fail to exhibit rapid equilibration in a particular style of non-CDA market, and their mode of failure is very similar to that of human traders in a similar non-CDA experiment reported by Smith [21].

Smith [21] reported results from an experimental model of common retail markets, where sellers announce prices and buyers either purchase at the offer price or ignore the offer, without giving any indication of what range of transaction prices they would be willing to bid. Smith's model was a modification of the CDA, rendered one-sided by preventing the buyers from quoting bid prices. Although this is a rather primitive approximation to retail markets (since superseded by experimental studies of *posted-offer* markets: see, e.g., [13, pp.173–239]), the results from Smith's experiment, and his explanation of those results, are intriguing. Smith's expectation was that transaction prices would settle at levels higher than the theoretical equilibrium price, indicating that the structure of retail markets offers advantages to the sellers. But this did not happen: instead, transaction prices settled at levels significantly below equilibrium. Smith explained this as being due to buyers that never quite recovered from having been 'badly fleeced' in the early stages of the experiment, where transactions occurred at high prices before equilibration had driven them lower.

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. In this paper, we illustrate these arguments by means of an example. We present results from experiments where ZIP traders adapt and interact via price-bargaining in market environments based on the artificial 'retail' market that Smith [21] used in his experimental economics research, and we demonstrate that groups of simple agents can exhibit human-like collective market behaviors. We note that, while it is often tempting to offer explanations of human market behavior in terms of the mental states of the agents in the market, our agents are sufficiently simple that mental states can have no useful role in explaining their activity. Thus, explanations of the human-like collective market behavior of our agents cannot be phrased in terms of mental states; thereby inviting comparisons with Braitenberg's influential "law of uphill analysis and downhill invention", with eliminative materialism in the philosophy of cognitive science, and with dynamical-systems-based analyses of adaptive behavior.

Section 2 introduces the mechanisms of adaptation in ZIP traders. In Section 3, we present results showing that ZIP-trader 'retail' markets exhibit the same failure qualities as Smith's human-trader 'retail' markets. From this, we argue that although the collective behavior of the trading animats is similar to that of the groups of humans, explanations of the animat markets could have significant impact on the way in which comparable human activity is explained. In particular, the simplicity of the ZIP trading mechanisms means that explanations of their failures cannot be phrased in terms of mental states such as not recovering from being fleeced during a trading day earlier in the experiment. We further discuss the implications of this in Section 4. In the remainder of this paper, we assume the reader is familiar with the basic microeconomics and details of experimental economics that we reviewed in the companion paper [9].

2 ZIP Traders

The emphasis in our work is on creating *simple* autonomous software agents, or animats, for bargaining in market-based environments. This emphasis on simplicity comes not only from a desire for computational efficiency (important in engineering applications if hundreds or thousands of animats are active on a network), but also in a speculative scientific attempt at sketching the minimum mechanistic complexity necessary and sufficient for explaining human bargaining behaviors in specific market environments.

Space restrictions prevent us from presenting a full discussion of the rationale for the current design of ZIP trader agents, and from presenting exhaustive results. The intention here is to briefly summarize key aspects of the design before presenting illustrative results. Cliff [6] gives a complete discussion of the design, shows results from many experiments in different types of market environment, and includes all the C source-code for the system. A recent thesis by van Montfort [25] replicated our CDA results, and explored the use of our ZI traders in spatially distributed markets where there may be potentially hundreds or thousands of traders.

In common with much work in (human-based) experimental economics, most of our studies to date have considered markets where each trading agent remains either a buyer or a seller for the duration of the entire experiment. However, van Montfort [25] demonstrated the use of our ZIP traders as arbitrage agents capable of buying units of commodity in one market for subsequent re-sale into another market, exploiting differences in price between the two markets.

Each ZIP trader operates by maintaining a *profit margin* that it uses for calculating the price it 'quotes' (offers or bids) in the market: the profit margin determines the difference between the price the agent quotes and the *limit price* for the commodity the agent is trading. For agents designated as sellers, the limit price is the price below which they may not sell a unit of the commodity. For agents designated as buyers, the limit price is the price above which they may not buy a unit of the commodity. Hence, when two traders enter into a transaction, the seller's profit is given by subtracting the seller's limit price from the transaction price, while the buyer's profit is given by subtracting the transaction price from the buyer's limit price.

The 'aim' of each ZIP agent is to maximize profit generated by trading in the market. If an agent's profit margin is set too low, it will miss out on potential profit when it makes a transaction with another agent, so all agents are constantly trying to increase their profit margins. But if an agent sets its profit margin too high, it may miss the opportunity to make transactions with other agents, because the price it offers is less attractive than the prices offered by competing agents. Clearly, what it means for the profit margin to be "too high" or "too low" is dependent on the context of the market conditions, and varies dynamically. Thus, the problem of designing a trading agent can be considered as a combination of two issues: the *qualitative* issue of deciding *when* to increase or decrease the profit margin, and the *quantitative* issue of deciding *by how much* the margin should be altered.

For reasons we discuss in detail in [6, 11], each ZIP trader makes the qualitative decision of when to alter its margin on the basis of four factors. The first factor is whether the agent is *active* in the market: agents are active until they have sold or bought their full entitlement of units of the commodity. The remaining three factors concern the last quote by any agent in the market: we refer to this as $\mathcal{Q}$. Each ZIP trader notes whether $\mathcal{Q}$ was an offer or a bid, whether $\mathcal{Q}$ was accepted (i.e., led to a transaction) or rejected (ignored by the traders in the market), and whether $\mathcal{Q}$'s price, $q(t)$, is greater than or less than the price the ZIP trader would currently quote. We refer to the price a ZIP trader i would quote at time t as that trader's *quote-price*, $p_i(t)$, which is calculated from i's limit price $\lambda_{i,j}$ (for i's jth unit of commodity) and i's current profit coefficient $\mu_i(t)$ using $p_i(t) = \lambda_{i,j}(1 + \mu_i(t))$. Thus, a seller's profit margin is raised by increasing μ_i and lowered by decreasing μ_i, such that $\mu_i(t) \in [0, \infty); \forall t \forall i$. The situation is reversed for buyers: they raise their margin by decreasing μ_i and lower it by increasing μ_i, subject to $\mu_i(t) \in [-1, 0]; \forall t \forall i$

A ZIP seller raises its profit margin whenever $\mathcal{Q}$ was accepted and $p_i(t) \leq q(t)$. It lowers its margin only if it is still active and $\mathcal{Q}$ was an offer with $p_i(t) \geq q(t)$, or if $\mathcal{Q}$ was a bid that was accepted and $p_i(t) \geq q(t)$. Similarly, a ZIP buyer raises its profit margin whenever $\mathcal{Q}$ was accepted and $p_i(t) \geq q(t)$, and it lowers its margin when it is active and either $\mathcal{Q}$ was a rejected bid with $p_i(t) \leq q(t)$ or $\mathcal{Q}$ was an accepted offer with $p_i(t) \leq q(t)$.

The quantitative issue of by how much the profit margin should be altered is addressed by using a simple machine-learning algorithm. Specifically, the learning rule we use is *Widrow-Hoff with momentum*, which also underlies back-propagation learning in neural networks [20]. Briefly, this adjusts the actual output of a system toward some *target* output value, at a speed determined by a learning rate β, and with a simple 'memory' or 'momentum' parameter γ. In each ZIP trader the target value $\tau_i(t)$ is given by a stochastic perturbation of $q(t)$, and each trader i uses this in combination with β_i and γ_i to adjust its profit-coefficient $\mu_i(t)$. The profit-margin update rule is:

$$\mu_i(t+1) = (p_i(t) + \Gamma_i(t))/\lambda_{i,j} - 1$$

where

$$\Gamma_i(t) = \gamma_i \Gamma_i(t-1) + (1-\gamma_i)\beta_i(\tau_i(t) - p_i(t))$$

and $\Gamma_i(0) = 0 : \forall i$.

The target price $\tau_i(t)$ is calculated by multiplying $q(t)$ by a *relative* coefficient $\mathcal{R}_i(t)$ and then adding a small *absolute* perturbation $\mathcal{A}_i(t)$. The values for $\mathcal{R}_i(t)$ and $\mathcal{A}_i(t)$ are stochastically generated from independent and identical distributions for each trader, every time $\tau_i(t)$ is calculated. When the trader's quote-price is being increased, $\mathcal{R}_i = \mathcal{U}(1.0, 1.0 + c_\mathcal{R})$ and $\mathcal{A}_i = \mathcal{U}(0.0, c_\mathcal{A})$, where $\mathcal{U}(c_{\text{lo}}, c_{\text{hi}})$ denotes a uniformly distributed random real value over the range $[c_{\text{lo}}, c_{\text{hi}}]$. When the trader's quote-price is being decreased, $\mathcal{R}_i = \mathcal{U}(1 - c_\mathcal{R}, 1.0)$ and $\mathcal{A}_i = \mathcal{U}(-c_\mathcal{A}, 0.0)$. For further details of how learning is implemented in ZIP traders, see Cliff and Bruten [6, 7].

In the experiments reported in this paper the following parameter values were used. Each trader's value for β_i was set randomly from $\mathcal{U}(\beta_{\text{lo}}, \beta_{\text{hi}})$ with β_{lo}=0.1 and β_{hi}=0.5. Each trader's value for γ_i was set randomly from $\mathcal{U}(\gamma_{\text{lo}}, \gamma_{\text{hi}})$ with γ_{lo}=0.0 and γ_{hi}=0.1. In generating $\tau_i(t)$, all traders use parameter-values $c_\mathcal{R} = 0.05$ and $c_\mathcal{A} = 0.05$. The initial profit coefficients (i.e., $\mu_i(0)$) of the traders were set randomly from uniform distributions symmetric about zero, determined by two parameters: $\mu_{\text{lo}} = 0.05$ and $\mu_{\text{hi}} = 0.35$: each seller's value of $\mu_i(0)$ was set randomly from $\mathcal{U}(\mu_{\text{lo}}, \mu_{\text{hi}})$ and each buyer's value of $\mu_i(0)$ was set randomly from $\mathcal{U}(-\mu_{\text{hi}}, -\mu_{\text{lo}})$.

All our work to date has involved experiments where the values of the system parameters have been determined manually (i.e., by trial and error). Our experience is that the system is fairly robust in the sense that it is not particularly sensitive to variations in the system parameters. Nevertheless, the use of some kind of automatic tuning or optimization technique such as a genetic algorithm is an obvious direction for possible future work: the free parameters in the current ZIP system are $c_\mathcal{A}$ and $c_\mathcal{R}$, and the pairs of upper and lower bounds (μ_{lo},μ_{hi}), ($\beta_{\text{lo}}, \beta_{\text{hi}}$), and ($\gamma_{\text{lo}}, \gamma_{\text{hi}}$); any or all of these could be placed under evolutionary control. Despite the documented difficulties of evolving situated autonomous agents for collective behaviors [17, 26], evolutionary optimization of adaptive trading agents offers the possibility of specializing generic adaptive traders to the structure and dynamics of particular markets.

3 One-Sided Auction 'Retail' Markets

In Smith's 1962 paper [21], all the experiments except one explored CDA markets. In the one non-CDA market, Smith examined the dynamics of a one-sided auction, where only sellers were allowed to quote offers: buyers were not allowed to quote bids, but passively observed the prices offered by the sellers. Each buyer therefore had the privilege of being able to ignore offer-prices that were too high and accept those that were within their range, without giving any indication of their limit prices. Smith proposed this as an approximation to an ordinary retail market, where sellers bear the responsibility of advertising their prices and buyers decide whether to buy or not without entering into any kind of bargaining or haggling process. Smith's results from this experiment are shown in Figure 1.

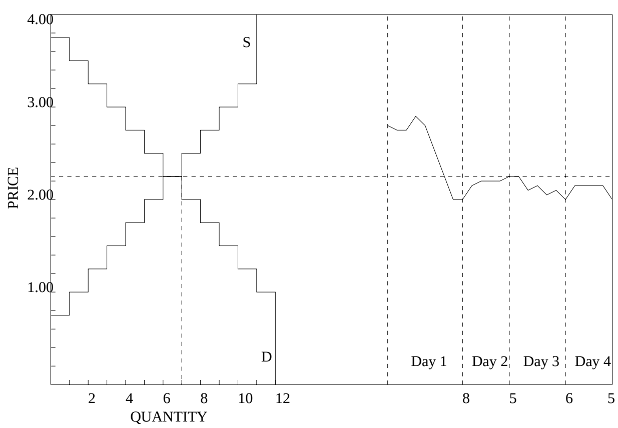

Figure 1: Left-hand side shows the market supply and demand, with theoretical equilibrium values P_0=\$2.25 and Q_0=7. Right-hand side shows time-series of transaction prices from four 'days' when traders interact via a one-sided (offer-only) 'retail' market. Numbers on the horizontal axis of theright-hand figure indicate market volume (quantity of transactions) for each day.

Smith's comments on his expectations and actual results for this experiment are significant:

> "Since sellers desire to sell at the highest prices they can get, one would expect the offer prices to be high, and, consequently, one might expect the exchange [i.e., transaction] prices to show a persistent tendency to remain above the predicted equilibrium. The result was in accordance with this crude expectation in the first market period [i.e., day] only.... Since sellers only were making offers, the prices tended to be very much above equilibrium. Five of these offers were accepted at prices ranging from \$2.69 to \$2.80... The competition of sellers pushed the offer prices lower and the remaining buyers made contracts at prices [of \$2.35, \$2.00, and \$2.00]. The early buyers in that first market period never quite recovered from having subsequently seen exchange prices fall much below the prices at which they had bought. Having been badly fleeced, through ignorance, in that first trading period, they refrained from accepting any high price offers in the remaining three periods of the test. This action, together with seller offer price competition, kept exchange prices at levels persistently below equilibrium for the remainder of [the experiment]." [21].

The ZIP traders can be used in a straightforward copy of Smith's experimental retail market. The supply and demand curves for the ZIP market are shown in Figure 2. For reference, Figure 3 shows the transaction-price time-series resulting from one experiment where the supply and demand curves shown in Figure 2 were used in a continuous double auction (CDA) market. As can be seen, the transaction prices of ZIP traders operating in a CDA market rapidly stabilize at values close to the theoretical equilibrium price of \$2.25. Figure 4 then shows the average results from 50 such experiments.

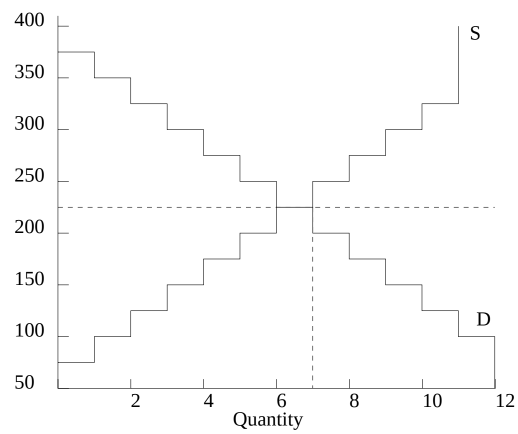

Figure 2: Supply and demand curves: 12 buyers and 11 sellers. Theoretical equilibrium price P_0 = \$2.25; quantity Q_0=7.

In Figure 5 we show the mean daily transaction prices from 50 experiments where the ZIP traders operate in Smith's 'retail market'. The same parameter values are used as in the experiments for Figures 3 and 4: the only difference is that the buyers are prevented from quoting bid-prices. As can be seen, the average transaction prices are typically less than \$2.00 (significantly below the theoretical equilibrium price of \$2.25). There also appears to be little or no convergence towards equilibrium, or reduction in variance as the experiment progresses. The apparent lack of convergence or reduction in variance can be better understood by examining individual price tra-

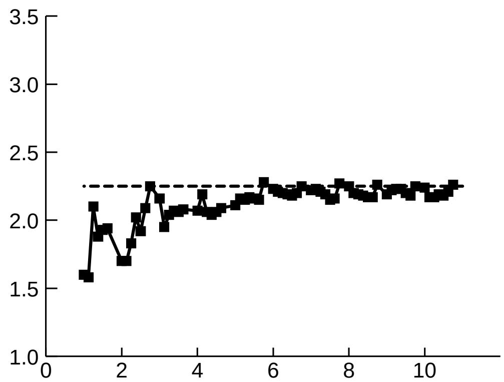

Figure 3: One transaction-price time-series from one experiment where the supply and demand of Figure 2 are used in a CDA market where both buyers and sellers can quote prices, for ten trading sessions or 'days'. The horizontal axis shows the day number, the vertical axis indicates the transaction price.

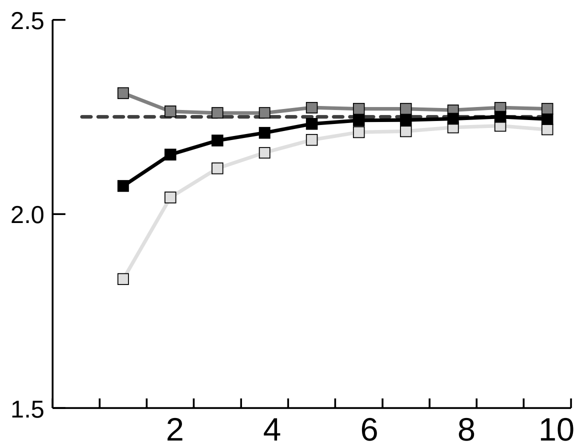

Figure 4: Mean transaction-price per trading session, averaged over 50 sets of results such as those shown in Figure 3. The horizontal dashed line shows the P_0 value. For each trading 'day', the graph shows the average value (black), and values plus (medium gray) and minus (light gray) one standard deviation, of the mean of the transaction prices in that day.

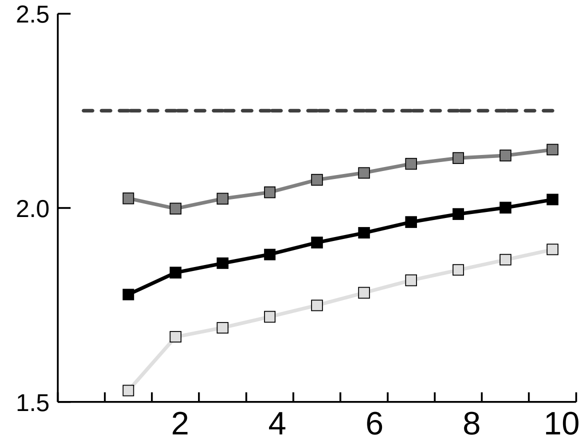

Figure 5: Mean ZIP transaction prices, averaged over 50 experiments, for 'retail-market' experiments with the supply and demand shown in Figure 2 (P_0=\$2.25). Format as for Figure 4

jectories: Figures 6 to 9 show time-series of the transaction prices in four individual experiments using ZIP traders in the 'retail' market with supply and demand as illustrated in Figure 2. As can be seen, in all four experiments the market converges to a fairly constant value for transaction prices by Day 4, but the value that is converged on varies: in Figures 6 to 8, all trades on Day 10 are within \$0.15 of the theoretical equilibrium, while in Figure 9 no trade is less than \$0.40 off the equilibrium price. As is clear in Figure 5, the price converged on is, on average, significantly less than the theoretical equilibrium. Thus, there is a strong qualitative agreement between results from our ZIP traders and Smith's [21] observations of human subjects in his experimental 'retail markets'.

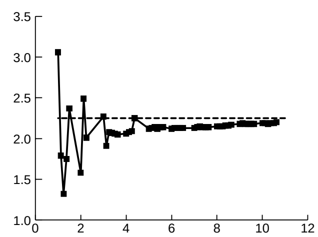

Figure 6: Transaction-price time series for one 'retail market' experiment, seed=3453.

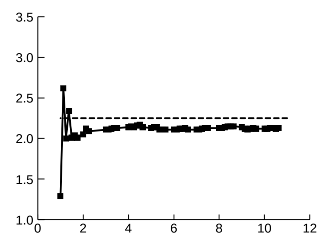

Figure 7: Transaction-price time series for one 'retail market' experiment, seed=3522.

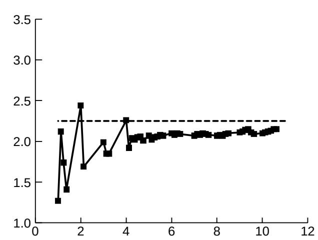

Figure 8: Transaction-price time series for one 'retail market' experiment, seed=3553.

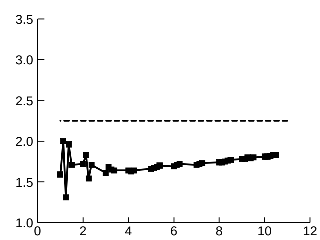

Figure 9: Transaction-price time series for one 'retail market' experiment, seed=3591.

Of these four single experiments, the price series in Figure 6 most closely resembles that of Smith's subjects: only three transactions occur at transaction price more than a few cents above the equilibrium price; while many more occur at prices lower than equilibrium, which is approached *very* slowly, from below. Smith's explanation was that this is due to early transactions at high prices preceding a series of low-price transactions that induce a resistance to higher prices in 'fleeced' traders. Whether this explanation can apply to our ZIP traders requires a more detailed examination of the dynamics of individual experiments. Figure 10 shows text output from Day 1 of the market experiment shown in Figure 6: in the first four transactions, sellers announce a price and one or more buyers are willing to buy at that price (the buyer who gets the deal is chosen at random from those that are willing). In the fifth transaction, Seller 10 makes an offer of \$3.52 which is ignored by the buyers: Seller 9 then offers at \$3.51; this is also ignored and Seller 9 offers again at \$3.50, which is again ignored; Seller 5 then offers at \$2.37, which is accepted by Buyer 0. For the sixth transaction, there is a sequence of 33 ignored offers, which ends when Seller 4 makes an offer of \$2.12 (having previously offered \$2.40, \$2.22, and \$2.16). For the seventh, there are 49 ignored offers before Seller 3 finally drops the offer price to \$2.07, and a deal is done. In the bargaining for the eighth transaction of the day, 100 quotes fail to find a taker, and the first day ends.

The effects this sequence of accepted and ignored offers has on the profit margins of the ZIP buyers and sellers is illustrated in Figure 11, which shows the apparent supply and demand curves and bid-and-offer arrays at the start of Day 1 and at the start of Day 2. As can be seen, the apparent supply and demand curves alter significantly over the first day. For intra-marginal units, the traders have increased their profit margins, flattening the supply and demand curves and bringing them closer together, thereby reducing the apparent surplus. For extra-marginal units, the traders have decreased their profit margins, again lessening the distance between the curves.

To better illustrate the alterations in the bid-and-offer arrays between the two states shown in Figure 11, Figure 12 shows the temporal progression of the arrays after each transaction in Day 1. As can be seen from the graphs labeled E to H, after four transactions the apparent supply and demand curves do not intersect, and so there is no theoretical equilibrium price or quantity. This gives rise to the sequences of ignored quotes illustrated in Figure 10 (5 before Figure 12E, 33 before Figure 12F, 49 before Figure 12G, and 100 before Figure 12H), which in turn lead to alteration of the traders' profit margins, thereby altering the apparent supply and demand so that eventually an intersection does occur, after which a transaction can take place. Typically, as soon as the apparent supply and demand curves intersect, two traders make a transaction and leave the market, and in doing so they alter the apparent supply and demand back to a state where no equilibrium is indicated.

```
day 1 trade 1
Seller  7 offers at 3.060  1 traders willing to deal
Seller  7 sells to Buyer 1

day 1 trade 2
Seller  2 offers at 1.790  5 traders willing to deal
Seller  2 sells to Buyer 3

day 1 trade 3
Seller  0 offers at 1.320  8 traders willing to deal
Seller  0 sells to Buyer 8

day 1 trade 4
Seller  1 offers at 1.750  6 traders willing to deal
Seller  1 sells to Buyer 6

day 1 trade 5
Seller 10 offers at 3.520  No willing takers (fails=1)
Seller  9 offers at 3.510  No willing takers (fails=2)
Seller  9 offers at 3.500  No willing takers (fails=3)
Seller  5 offers at 2.370  1 traders willing to deal
Seller  5 sells to Buyer 0

day 1 trade 6
Seller  3 offers at 2.210  No willing takers (fails=1)
Seller  6 offers at 2.520  No willing takers (fails=2)
Seller  6 offers at 2.530  No willing takers (fails=3)
Seller  8 offers at 2.930  No willing takers (fails=4)
Seller 10 offers at 3.350  No willing takers (fails=5)
Seller  9 offers at 3.080  No willing takers (fails=6)
Seller  8 offers at 2.820  No willing takers (fails=7)
Seller  9 offers at 3.040  No willing takers (fails=8)
Seller  9 offers at 3.010  No willing takers (fails=9)
Seller  4 offers at 2.400  No willing takers (fails=10)
...
Seller  4 offers at 2.220  No willing takers (fails=15)
...
Seller  4 offers at 2.160  No willing takers (fails=24)
...
Seller  8 offers at 2.780  No willing takers (fails=32)
Seller  9 offers at 3.010  No willing takers (fails=33)
Seller  4 offers at 2.120  1 traders willing to deal
Seller  4 sells to Buyer 4

day 1 trade 7
Seller 10 offers at 3.350  No willing takers (fails=1)
...
Seller  8 offers at 2.780  No willing takers (fails=49)
Seller  3 offers at 2.070  1 traders willing to deal
Seller  3 sells to Buyer 2 1.180)

day 1 trade 8
...
Seller  8 offers at 2.760  No willing takers (fails=100)
```

Figure 10: Text output showing quotes and transactions for Day 1 in the experiment of Figure 6. Much text has been deleted to increase clarity.

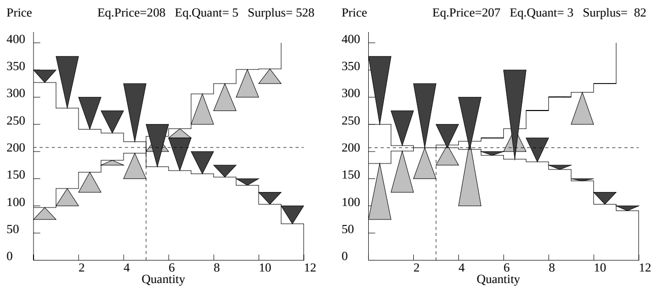

Figure 11: Bid-and-offer arrays in the experiment of Figure 6. Left: at the start of Day 1. Right: at the start of Day 2. Limit and quote prices are indicated using the format introduced in [9]: Each buyer's limit and quote prices are illustrated as dark inverted triangles, while each seller's limit 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.

Figure 13 shows the bid-and-offer arrays at the start of each subsequent day in the experiment. As is clear, although the rank ordering of the traders varies as they alter their prices up or down by a few cents, there is very little change in the overall shape of the bid-and-offer arrays after Day 3. The fact that in this experiment the market converges on transactions around \$2.12 (i.e., less than the theoretical equilibrium price of \$2.25) is consistent with Smith's [21] results from his experiment with human subjects, where transaction prices also converged to a stable below-equilibrium level.

Thus, in addition to our demonstration in other publications [6, 7, 10, 11] that ZIP traders can give human-like collective behavior in CDA markets, the results presented here show that the dynamics and the modes of failure of ZIP traders are also similar to those of humans in Smith's [21] one-sided auction experimental model of retail markets. The implications of this are discussed further in the next section.

To demonstrate that the difference in market organization (i.e., the difference between the CDA and one-sided 'retail' auction rules) accounts for the differences seen in the transaction-price data of the two markets, we close this section with the data in Figures 14 and 15. Both of these figures show price data from markets where the market organization is 'retail' for the first five days and then switches to CDA for the remaining ten days. As can be seen, once the market alters from retail to CDA, the transaction prices of the ZIP traders rapidly approaches the theoretical competitive equilibrium. Note that the *only* change is in the market organization: all other parameters remain the same, and none of the trader's variables (e.g. $\mu_i(t)$ or $\Gamma_i(t)$) are altered when the organization is changed. Clearly then, the market organization is a primary cause of the equilibration failure.

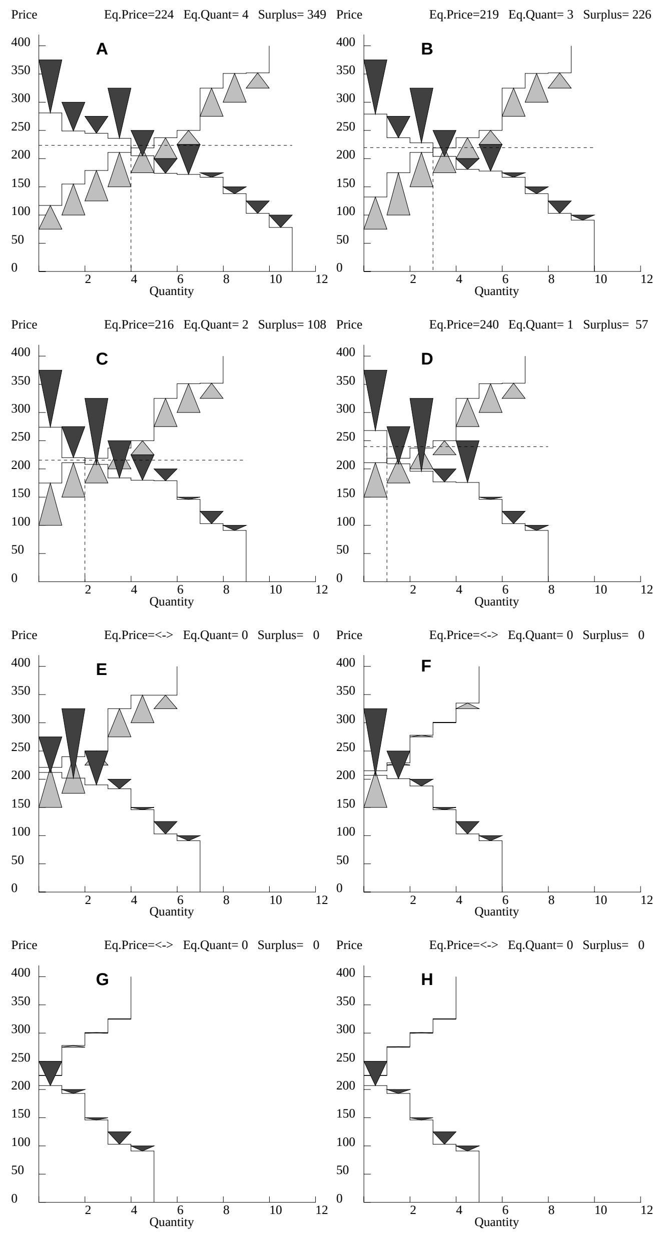

Figure 12: Temporal progression of bid-and-offer arrays for days 1 to 2 in the price series shown in Figure 6. Each graph shows the bid-and-offer arrays of the active traders after a transaction: A is after the first transaction; B is after the second transaction; And so on until H which is after the eighth (end of Day 1).

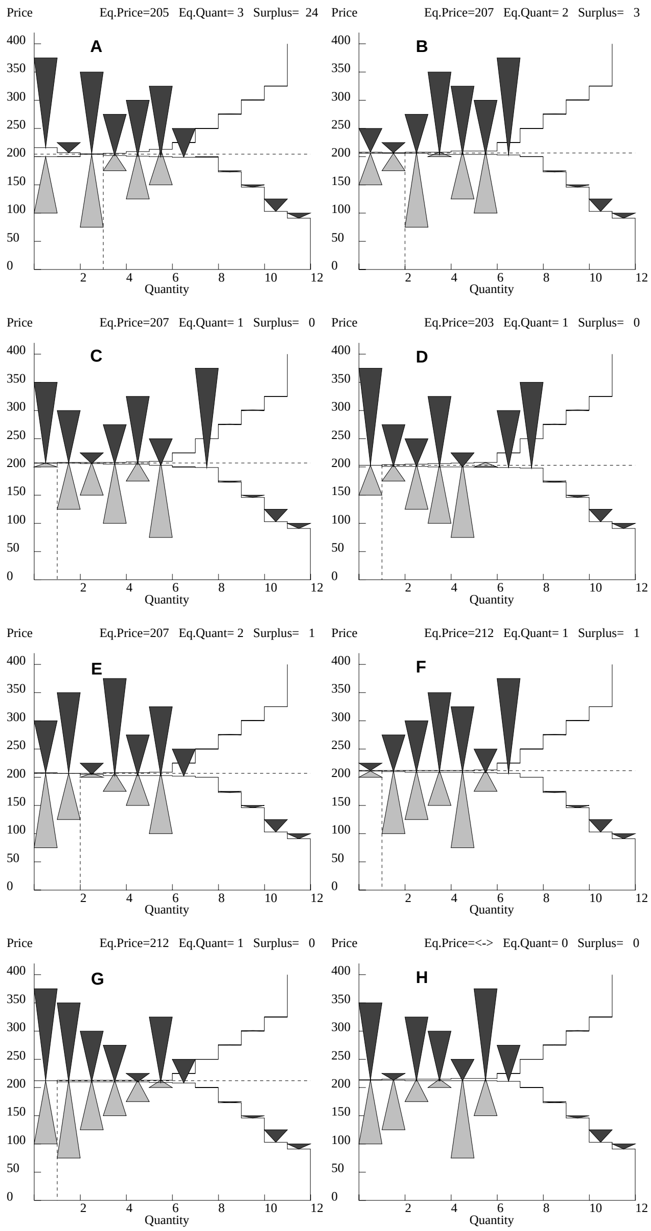

Figure 13: Temporal progression of bid-and-offer arrays for days 3 to 10 in the price series shown in Figure 6. Each graph shows the bid-and-offer array at the start of a day's trading: A is day 3; B is day 4; and so on until H which shows the start of day 10. See text for discussion.

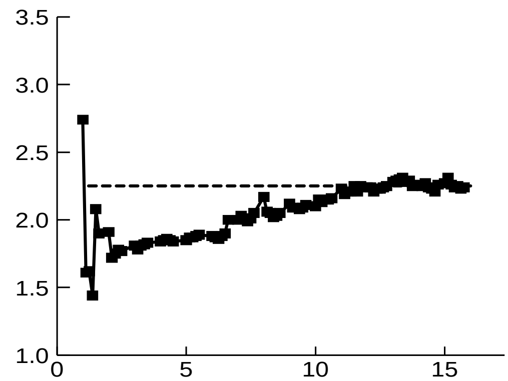

Figure 14: One transaction-price time-series from one experiment where the supply and demand of Figure 2 are used in a 'retail' market for the first 5 days, before switching to a CDA market for the last 10 days. The horizontal axis shows the day number, the vertical axis indicates the transaction price.

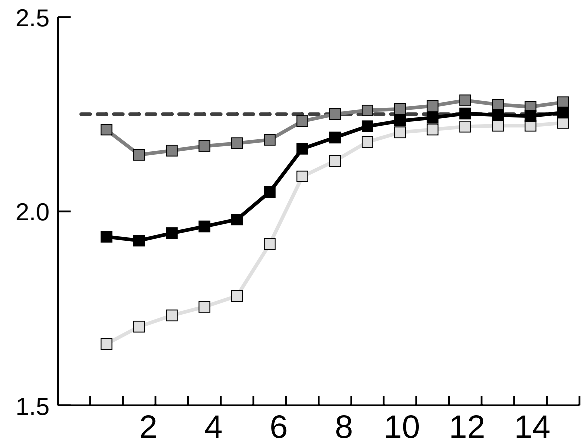

Figure 15: Mean transaction-price per trading session, averaged over 50 sets of results such as those shown in Figure 14.

Finally, it is sobering to note that with synthetic adaptive agents it is possible to record all manner of significant variables, both internal and external to the agent, and to visualize them in styles such as those shown in Figures 5 to 13. And this is from just one experiment, which took less than five seconds to run on a medium-power workstation (a Sun Sparc20). Clearly, hundreds or thousands of experiments can be run with artificial agents in the time it takes one experiment to be conducted with human subjects. Indeed, with one workstation, in one week it would be possible to run approximately 100,000 artificial-agent experiments: this is probably more experiments than have been run with human subjects in the entire history of experimental economics. But this is not necessarily an advantage: each experiment has the potential to generate masses of data; managing, visualizing, and analyzing the data to arrive at meaningful conclusions could present serious problems, and should be noted as a topic for further work.

4 Discussion

The similarity between our ZIP results and those from Smith's human subjects suggests a line of reasoning similar to that underlying much adaptive behavior research. This reasoning relies on noting that there is one key difference between our results and Smith's. Smith was working with human subjects, where there is a natural temptation to offer explanations in terms of mental states. In the passage quoted above, Smith talks of the human buyers "never quite recovering" from "having been badly fleeced". It is not clear from the original text whether this account is inventive conjecture on Smith's part, or the result of properly conducted post-experiment interviews. But even if these comments are the result of interviewing those subjects who ended up as 'fleeced' buyers, the danger of introspective *a posteriori* accounts of behavior are well known.

The crucial difference then, between Smith's work with humans and our work with ZIP traders, is that in the ZIP traders there are no place for such mentalistic descriptions of the behavior of the agents in the market. There is nothing, not even an evolved neural network, in which the ZIP agents could hide the mental states of 'never quite recovering' or noticing that they have been 'badly fleeced'. Any explanation of what causes the ZIP-agent markets to approach equilibrium slowly and from below is *forced* to be framed in terms of the interactions of the simple ZIP adaptation mechanisms, because there is *nothing else* in the system that could cause the observable phenomena.[1] Let us assume that a causal mechanistic explanation for how the collective behavior of ZIP traders gives rise to some market-level phenomena can be developed, and call it $\mathcal{E}$. Then $\mathcal{E}$ can also be considered a candidate explanation for the behavior of groups of human traders. Naturally, if it can be demonstrated that the ZIP traders are using adaptation mechanisms that could not be employed or implemented by humans, then $\mathcal{E}$ is a very weak explanation, or no explanation at all. But if such counter-arguments to $\mathcal{E}$ cannot be readily advanced, $\mathcal{E}$ should properly be considered as a putative explanation for the human behavior, which can be subjected to experimental evaluation or falsification. And, crucially, $\mathcal{E}$ cannot be phrased in terms of mental or emotional states, because the ZIP traders have nothing that corresponds to such states.

The failure of ZIP traders to converge on a competitive equilibrium (i.e., a steady sequence of transaction prices at the P_0 value) in 'retail' markets is due simply to the fact that although the buyer and seller profit-margins are altered symmetrically when in a CDA, the prevention of bids in the one-sided 'retail' market introduces an asymmetry: although the traders raise their margins under symmetric conditions, an active buyer b lowers its margin only when $\mathcal{Q}$ was *accepted* at a price $q(t) \geq p_b(t)$, while an active seller s will lower its margin when $q(t) \leq p_s(t)$ regardless of whether $\mathcal{Q}$ was accepted or not. In essence, this demonstrates that, despite the good equilibration properties of CDA markets where both buyers and sellers are trading according to the ZIP strategy described in Section 2, the asymmetry of opportunity sets (i.e., the prevention of bids) in the 'retail' market prevents equilibration by ZIP traders because their trading strategies depend on the bilateral flow of information found in CDA markets. While it may be possible to alter the ZIP strategies to give good equilibration in retail markets, or even in both retail markets *and* CDA markets, the key issue here is that our explanation of ZIP traders' failure to reach a competitive equilibrium is not reliant on them having vague and difficult-to-define mental states such as 'never quite recovering from being badly fleeced'.

By specifying and observing simple synthetic trading agents, it is possible to demonstrate the same overall market behavior without relying on abstract or vague descriptions of the mental states of the participants in the market. In this sense, the work described here is similar to other work in adaptive behavior that is justified by the principle that it can be more fruitful and more parsimonious to attempt an understanding of how some behavior is generated by *synthesising* an artificial system that exhibits that behavior, rather than by *analyzing* a natural system that exhibits the same behavior: a principle that Braitenberg [2] named the "law of uphill analysis and downhill invention". Although it is often difficult to resist the temptation to describe the cognitive behaviors of animals (and humans in particular) in terms of mental states, there is growing support for a counter-approach, where the intention is to explain observed behaviors in terms of the dynamics of causal mechanistic interactions, rendering the mental-states-based accounts obsolete. These ideas first gained credence in the philosophy of mind, where they are most strongly associated with Churchland's *eliminative materialism* [3, 4], and their relevance to work in artificial autonomous agents has been discussed by Smithers [22], van Gelder & Port [23, 19, 24], and Cliff & Noble [12]. Thus, our work here can be viewed as a step in the direction of adopting an eliminative materialism or dynamical systems perspective on human economic activity.

So, we have demonstrated here that ZIP traders can give results qualitatively similar to those of humans in 'retail market' experiments. In doing so, we have demonstrated a point of more general significance: that techniques common in adaptive behavior research can be

[1] Assuming, of course, that the code for the system has no bugs.

used to cast new lines of inquiry on the human experimental economics data. Given that ZIP traders exhibit human-like behavior and have no mental states, of how much genuine use are mental states in the explanation of human market behavior?

5 Conclusion

The development of computational mechanisms that allow groups of software agents to exhibit bargaining behaviors in market-based environments satisfies a number of needs. In market-based control, simple mechanisms are required to give computationally efficient, robust, and truly distributed resource allocation and control. Such mechanisms could also be employed in the growing field of internet-based commerce. Moreover, such mechanisms act as mechanistically rigorous statements of potential models of human bargaining behaviors, although it is likely that more complex mechanisms would be required to further account for the many subtleties and nuances of human behavior: empirical work in experimental economics and human psychology would also be necessary to validate any models. Once validated, such model agents could be used in the manner intended in the work of Arthur [1] or Easley and Ledyard [14], for conveniently testing theories concerning the behavior of humans in different market structures and conditions.

The arguments we presented in earlier papers [6, 8, 9] indicate a need for bargaining mechanisms more complex than the constrained stochastic generation of bid and offer prices used by Gode and Sunder's "zero-intelligence" (ZI) traders [15]. The work on ZIP traders, reported here and in other papers [6, 7, 10] should be viewed as a preliminary sketch of what forms such bargaining mechanisms might take. The ZIP traders are more complex than Gode and Sunder's ZI traders, but only slightly, and in any case are manifestly much less complex than humans. Nevertheless, the results from the ZIP traders, both in terms of equilibration in CDA markets and failure to equilibrate in Smith's one-sided auction model of retail markets, are clearly closer to those from human experimental markets than are the results from ZI traders. It is reassuring to see that the ZIP mechanisms can give such human-like results, but there is much further work that could be done in exploring behavior of ZIP traders in more complex market environments, and in attempting to extend the behavioral sophistication of such traders without unduly adding to their complexity.

References

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

[2] V. Braitenberg. *Vehicles: Experiments in Synthetic Psychology*. MIT Press, 1984.

[3] P. M. Churchland. *Scientific Realism and the Plasticity of Mind*. Cambridge University Press, 1979.

[4] P. M. Churchland. *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*. MIT Press, 1989.

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

[6] D. Cliff. Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical Report HPL–97–91, Hewlett Packard Labs, 1997.

[7] D. Cliff & J. Bruten. Less than human: Simple adaptive trading agents for CDA markets. Technical Report HPL–97–155, Hewlett Packard Labs, 1997.

[8] D. Cliff & J. Bruten. More than zero intelligence needed for continuous double-auction markets. Technical Report HPL–97–157, Hewlett Packard Labs, 1997.

[9] D. Cliff & J. Bruten. Shop 'til you drop I: Market trading interactions as adaptive behavior. Technical Report HPL–97–1nn, Hewlett Packard Labs, 1997.

[10] D. Cliff & J. Bruten. Simple bargaining agents for decentralized market-based control. Technical Report HPL–98–17, Hewlett Packard Labs, 1997.

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

[12] D. Cliff & J. Noble. Knowledge-based vision & simple visual machines. *Phil. Trans. Roy. Soc. Lond.: B*, 352:1165–1175, 1997.

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

[14] D. Easley & J. Ledyard. Theories of price formation and exchange in double oral auctions. In D. Friedman & J. Rust, eds, *The Double Auction Market: Institutions, Theories, and Evidence*, pp.63–98. Addison-Wesley, 1992.

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

[16] J. Hagel & A. Armstrong. *Net Gain: Expanding Markets Through Virtual Communities*. Harvard Business School Press, 1997.

[17] M. Matarić & D. Cliff. Challenges in evolving controllers for physical robots. *Robotics & Auton. Sys.*, 19(1):67–83, 1996.

[18] L. W. McKnight & J. P. Bailey, eds. *Internet Economics*. MIT Press, 1997.

[19] R. Port & T. van Gelder, eds. *Mind as Motion: Explorations in the Dynamics of Cognition*. MIT Press, 1995.

[20] D. Rumelhart, G. Hinton, & R. Williams. Learning internal representations by error propagation. In D. Rumelhart & J. McClelland, eds, *Parallel Distributed Processing, Volume 1: Foundations*, pp.318–362. MIT Press, 1986.

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

[22] T. Smithers. Taking eliminative materialism seriously: A methodology for autonomous systems research. In F. J. Varela & P. Bourgine, eds, *Towards a Practice of Autonomous Systems: Proc. ECAL91*, pp.31–40. MIT Press, 1992.

[23] T. van Gelder. What might cognition be if not computation? Technical Report 75, Indiana Uni. Cognitive Sciences, 1992.

[24] T. van Gelder. The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, forthcoming, 1998.

[25] G. van Montfort. Economic agents for controlling complex systems. Master's thesis, Delft University of Technology, September 1997. Unpublished.

[26] N. Zaera, D. Cliff, & J. L. Bruten. (Not) evolving collective behaviors in synthetic fish. In P. Maes, M. J. Matarić, J.-A. Meyer, J. Pollack, & S. W. Wilson, eds, *From Animals to Animats 4: Proc. SAB96*, pp.635–644, MIT Press, 1996.