Less Than Human: Simple adaptive trading agents for CDA 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

Gode and Sunder (1993) described continuous double-auction (CDA) markets populated by "zero-intelligence" (ZI) traders that act randomly. Their results appear to indicate that no intelligence is required to give human-like trading performance. Cliff and Bruten (1997) demonstrated serious failings of ZI traders. Here, 'ZI-plus' (ZIP) traders are introduced: simple stochastic agents that adapt over time using an elementary form of machine learning. It is shown that the performance of ZIP traders is significantly closer to the human data than is that of ZI traders. Thus, human-like trading behavior can be achieved with intelligence that is more than zero but much less than human.

1 INTRODUCTION

Gode and Sunder (1993) reported results from market experiments where "zero-intelligence" (ZI) programs that submit random bids and offers replace human traders in experimental continuous double-auction (CDA) markets. Gode and Sunder (G&S herein) demonstrated that the transaction-price time-series of "ZI-constrained" (ZI-C) traders were *human-like* insofar as they appeared to converge on the theoretical equilibrium price and yielded levels of allocative efficiency similar to those of comparable human markets. G&S concluded that allocative efficiency was a poor indicator of the intelligence of agents in a CDA market, and noted that ZI-C traders could be distinguished from humans in CDA markets because the ZI-C traders gave higher levels of profit dispersion.

Cliff and Bruten (1997) analyzed the probability functions underlying ZI-C traders' interactions, in a variety of CDA markets, and predicted market conditions in which ZI-C traders would fail to trade at equilibrium prices. Empirical results from experiments demonstrated ZI-C traders failing as predicted. Thus, Cliff and Bruten (1997) claim that the ZI-C traders lack sufficient intelligence or rationality to exhibit human-like equilibration in CDA markets. On face value, G&S's ZI traders could be assumed to represent a lower limit on the intelligence requirements of traders in CDA markets, but Cliff & Bruten's results indicate that such an assumption would be unfounded.

This paper introduces simple adaptive trading agents for CDA markets, referred to as "zero-intelligence-plus" (ZIP) traders, that employ an elementary form of 'learning' to adapt their trading behavior on the basis of past experience of the market. It is shown that the ZIP traders avoid the failures of the ZI-C traders, and give human-like market behavior despite having intelligence that is much less than human.

2 ZERO INTELLIGENCE TRADERS

2.1 Background

Smith (1962) demonstrated that the transaction prices of remarkably small groups of human traders, operating in experimental CDA markets, rapidly approach the theoretical equilibrium price. But the question of just how much intelligence is required of an agent to achieve human-level trading performance is an intriguing one. This question was addressed by G&S (1993), whose results appear to indicate that no intelligence at all is required of the traders, so long as they are prevented from trading at a loss.

G&S used ZI-C trader-programs in CDA markets. They found that the imposition of a budget constraint (that prevents ZI-C traders from entering into loss-making deals), is sufficient to raise the allocative efficiency of the auctions to values near 100%. They conclude that the traders' motivation, intelligence, or learning have little effect on the allocative efficiency, which derives instead largely from the structure of the CDA markets. Thus, they claim, "Adam Smith's invisible hand may be more powerful than some may have thought; it can generate aggregate rationality not only from individual rationality but also from individual irrationality." (1993, p.119).

This important work has been influential in the experimental economics literature: Cliff and Bruten (1997) supply a list of 12 example publications that approvingly cite G&S's result.

G&S used an electronic CDA market, where traders were connected on a computer network. G&S's experiments with human traders were performed in a manner similar to that established by Smith (1962): the subjects were divided into a group of sellers and a group of buyers. Sellers were given a number of units of an arbitrary commodity, and each unit had a limit price (below which it could not be sold), which was private (i.e., known only to the seller of that unit). Buyers were given the rights and means to buy a number of units, and for each unit they were given a private limit price, above which they could not pay. The array of sellers' limit prices determined the market supply curve, and the array of buyers' limit prices determines the market demand curve. In the experiments with human traders, traders 'quote' bid and offer prices by typing them into their computer terminals: the quotes were then distributed to the other traders, and at any time a buyer could accept a seller's offer or a seller could accept a buyer's bid. This continuous trading process was broken into discrete periods or 'days': at the start of each day, new allocations of selling or buying rights were distributed to the traders. In experimental CDA markets such as these, as with real human CDA markets, transaction prices rapidly approached the theoretical equilibrium value given by the intersection of the supply and demand curves.

G&S replaced the humans with ZI-C traders, simple programs or "software agents". G&S used ZI-C traders in markets with supply and demand curves similar or identical to those used with their human subjects. Each ZI-C trader generates random bid or offer prices, but using a distribution *constrained* by the limit price for the current unit: each buyer is constrained to bid a price chosen randomly in the range $[1, \lambda_b]$ where λ_b is that buyer's limit price; each seller is constrained to offer at a price chosen randomly from the range $[\lambda_s, 200]$ where λ_s is that seller's limit price.

G&S presented results from five types of market. For each type of market, they showed time-series of transaction prices from one experiment with ZI-C traders, and from one experiment with human traders. The surprising and significant observation that G&S make is that the results from ZI-C traders appear to be very similar to those of human traders. In particular, G&S monitored allocative efficiency (profit extracted from the market as a proportion of maximum possible profit in that market) and found that the allocative efficiency of humans and ZI-C traders were not significantly different. Thus, they conclude that no intelligence other than the budget constraint is required of trading agents to exhibit human-like behavior in CDA markets. G&S speculate that no intelligence is necessary for the transaction prices of the traders to converge to the equilibrium value, and they close their paper with brief discussion of measurements of profit dispersion, $\mathcal{D}$, defined for a group of n traders as $\mathcal{D} = (\frac{1}{n}\sum_{i=1}^{n}(a_i - e_i)^2)^{0.5}$. where a_i is the actual profit earned by trader i, and e_i is the profit that trader would realize if all units were traded at the equilibrium price. Values of $\mathcal{D}$ for ZI-C traders are significantly higher than those for human CDA markets.

2.2 Critique

Cliff and Bruten (1997) presented a critique of G&S's (1993) work, motivated by considering ZI-C trading behavior in four types of market, three of which differ from those used by G&S. Fig. 1 shows

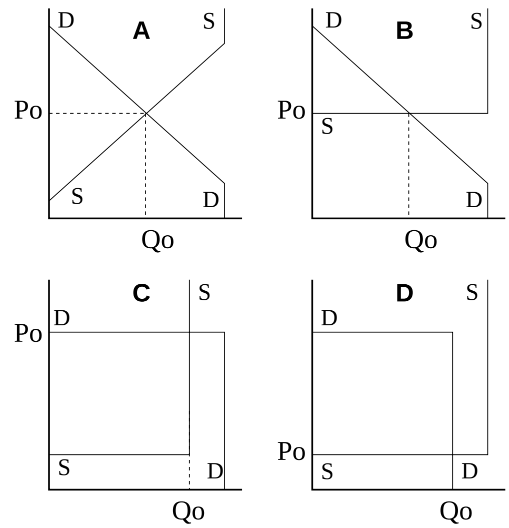

Figure 1: Four types of market. In each graph, the horizontal axis is quantity and the vertical is price. The supply curve is labeled SS and the demand curve is labeled DD; the intersection of these curves gives the equilibrium price P_0 and quantity Q_0. At top left is a market labeled A, where the supply and demand curves are symmetric about the line of constant price P_0. The top right market, labeled B, has a flat supply curve. At bottom left there is a 'box' market with excess demand, labeled C. At bottom right there is an excess-supply 'box' market labeled D.

the supply and demand curves for the four types of market, labeled A, B, C, and D, with equilibrium price P_0 and quantity Q_0 as indicated.

In market A, the supply and demand curves are *symmetric*, in that they have gradients that are approximately equal in magnitude but opposite in sign. In Market B, the supply curve is flat over the range of quantities supplied. In Markets C and D, both the supply curve *and* the demand curve are flat. However, in C, demand exceeds supply, and so the equilibrium price P_0 is set by the point where the supply curve cuts up through the demand curve, because the excess demand encourages price competition among buyers that will lead to bid-price increases until the maximum buyer limit price is reached. Similarly, in D, supply exceeds demand and so the excess supply encourages offer-price cuts, driving the price down to equilibrium at the point where the demand curve cuts down through the supply curve.

In the five experiments presented by G&S, the market supply and demand were all similar to A, although not so perfectly symmetric over the range of quantities 0 to Q_0. Yet markets such as as B, C, and D have also been studied in the literature. For example, market B is similar to Smith's (1962) "Chart 4", and markets C and D are similar to Smith's (1962) "Chart 6". Markets C and D are also known as "box-design" schedules (Davis & Holt, 1993, p.141).

For each of the four styles of market shown in Fig. 1, Cliff and Bruten (1997) developed analytic predictions of how the ZI-C traders would perform: in markets B, C, and D, failure of the average ZI-C transaction prices to reach equilibrium was predicted from analysis and confirmed empirically.

3 ZI-PLUS TRADERS

It is demonstrated here that remarkably simple adaptive mechanisms can give performance that does not suffer from the problems affecting G&S's ZI-C traders: only a slight increase in 'intelligence' is necessary. Thus, these trading agents are referred to as "zero-intelligence-plus" (ZIP) traders.

A profit-motivated trader in an experimental CDA market will not quote an initial offer-price or bid-price for a unit that is equal to that unit's limit-price λ: to do so would be to trade at zero profit. If p denotes the price a trader quotes for a transaction, then sellers should quote $p > \lambda$ and buyers should quote $p < \lambda$. The relative difference between p and λ will be referred to as the trader's *profit margin.* The ZIP traders adjust their profit margins up or down, on the basis of the prices of bids and offers made by the other traders, and whether those quotes are accepted (leading to transactions), or ignored.

The intention here is only to demonstrate that the simple adaptive mechanisms in ZIP traders can give results better than ZI-C traders and more similar to those of human traders. The complete rationale for the current design of ZIP trader agents, and extensive sets of results, are given in Cliff (1997), which includes all the C source-code for the system programs. A recent thesis by van Montfort (1997) explores the use of these ZIP traders in spatially distributed markets where there may be hundreds or thousands of agents.

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 (discussed in Section 3.1), and the *quantitative* issue of deciding *by how much* the margin should be altered (discussed in Section 3.2). Having explained these details, Section 3.3 then presents results from experiments with ZIP traders operating in the markets used to illustrate the failure of ZI-C traders.

3.1 Qualitative Considerations

To eliminate the need for sophisticated memory mechanisms, each ZIP trader alters its profit margin on the basis of four factors. The first factor is whether the agent is *active* in the market: agents are active if they are still capable of making a transaction in the market, and are inactive if they have sold or bought their full entitlement of units of the commodity. The remaining three factors concern q, the last price quoted by any agent in the market: referred to as the *last quote.* Each ZIP trader takes note of whether the last quote was an offer or a bid, whether it was accepted (i.e., led to a transaction) or rejected (ignored by the traders in the market), and whether it is greater than or less than the price that ZIP trader would currently quote. The price a ZIP trader i would currently quote is trader i's *quote-price*, p_i, which is calculated from trader i's values of $\lambda_{i,j}$ (i's limit price for unit j) and μ_i (i's profit margin) using $p_i = \lambda_{i,j}(1 + \mu_i)$.

A ZIP seller raises its profit margin whenever the last quote was accepted and $p_i \leq q$. It lowers its margin only if it is still active and the last quote was an offer with $p_i \geq q$, or if the last quote was a bid that was accepted and $p_i \geq q$. Similarly, a ZIP buyer raises its profit margin whenever the last quote was accepted and $p_i \geq q$, and it lowers its margin when it is active and either the last quote was a rejected bid with $p_i \leq q$ or the last quote was an accepted offer with $p_i \leq q$. The adaptation mechanism that alters the profit margin is described next.

3.2 Adaptation

At a given time-step t, a ZIP trader i calculates its quote-price $p_i(t)$ for its unit j with limit price $\lambda_{i,j}$ using the current value of its profit margin $\mu_i(t)$, thus: $p_i(t) = \lambda_{i,j}(1 + \mu_i(t))$, where $\mu_i(t) \geq 0.0$ for sellers and $\mu_i(t) \in [-1, 0]$ for buyers. The ZIP traders alter their profit margins using a simple update rule: $\mu_i(t+1) = (p_i(t) + \,,_i(t))/\lambda_{i,j} - 1$, where $,_i(t) = \gamma_i \,,_i(t-1) + (1-\gamma_i)\Delta_i(t-1)$ is the amount of change on the transition from t to $t+1$. Here, γ_i is the *momentum coefficient* for trader i, and trader i's value for $\Delta_i(t)$ is given by:

$$\Delta_i(t) = \beta_i(\tau_i(t) - p_i(t)) \quad (1)$$

Where β_i is the *learning rate* for trader i, and $\tau_i(t) = \mathcal{R}_i(t)q(t) + \mathcal{A}_i(t)$ is the *target price*, with $q(t)$ the price-value of the last quote made by a trader in the market, and $\mathcal{R}_i(t)$ and $\mathcal{A}_i(t)$ are stochastic values discussed further below.

The core of this adaptation method is Equation 1, the Widrow-Hoff "delta rule" which also underlies learning in back-propagation artificial neural networks (Rumelhart, Hinton, & Williams, 1986) and classifier systems (Wilson, 1994). This rule gives asymptotic convergence of $p_i(t)$ to a target price $\tau_i(t)$ at a speed determined by $\beta_i \in [0, 1]$.

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 intention is to increase the trader's quote price, $\mathcal{R}_i = \mathcal{U}(1.0, 1.0 + c_{\mathcal{R}})$ and $\mathcal{A}_i = \mathcal{U}(0.0, c_{\mathcal{A}})$, where $\mathcal{U}(c_l, c_h)$ denotes a uniformly distributed random real value over the range $[c_l, c_h]$. When the intention is to decrease the trader's quote-price, $\mathcal{R}_i = \mathcal{U}(1 - c_{\mathcal{R}}, 1.0)$ and $\mathcal{A}_i = \mathcal{U}(-c_{\mathcal{A}}, 0.0)$.

d	A:ZI-C	A:ZIP	B:ZI-C	B:ZIP
1	201 ± 12	185 ± 19	233 ± 09	225 ± 07
2	201 ± 14	191 ± 11	234 ± 10	208 ± 05
4	199 ± 14	196 ± 06	236 ± 10	201 ± 01
6	199 ± 13	199 ± 04	232 ± 10	201 ± 01
10	201 ± 15	200 ± 02	235 ± 10	201 ± 01

Table 1: Results for ZI-C and ZIP traders in markets A and B: $P_0 = 200$. Each column shows the mean ($\pm$ standard deviation) transaction price for day d.

d	C:ZI-C	C:ZIP	D:ZI-C	D:ZIP
1	138 ± 14	126 ± 20	249 ± 14	240 ± 08
2	137 ± 17	129 ± 17	246 ± 13	237 ± 09
4	135 ± 17	138 ± 19	253 ± 12	228 ± 10
6	136 ± 16	143 ± 20	251 ± 12	222 ± 10
10	137 ± 15	152 ± 21	248 ± 13	213 ± 11
30	136 ± 16	186 ± 20	250 ± 12	203 ± 07

Table 2: Results for ZI-C and ZIP traders in markets C and D, format as for Table 1.

The Widrow-Hoff rule forces $p_i(t)$ to approach $\tau_i(t)$, but $\tau_i(t)$ is itself dynamically varying. The system may be 'damped' by setting $\gamma_i \in [0,1]$ to non-zero values, limiting high-frequency oscillations. This method is also used in back-propagation neural networks (Rumelhart et al., 1986).

Data from markets of ZIP traders with this profit-margin adaptation are shown below.

3.3 Results

To allow direct comparison, results are presented here from ZIP traders operating in the four markets that Cliff and Bruten (1997) used to show the failure of ZI-C traders: the supply and demand curves for these markets are as illustrated in Fig.1. In all markets, the equilibrium values were $P_0 = 200$ and $Q_0 = 6$. To give a representative view of the performance of ZIP traders, all experiments presented here are conducted with the same parameter values, rather than with values optimized or 'tuned' to give good performance for each particular market.

When initializing each trader, β_i is generated from $\mathcal{U}(0.1, 0.5)$, and γ_i is generated from $\mathcal{U}(0.0, 0.1)$. Both β_i and γ_i remain fixed for the duration of the experiment. Initial values for the μ_i profit margins of the traders are generated using $\mathcal{U}(0.05, 0.35)$ for sellers and $\mathcal{U}(-0.35, -0.05)$ for buyers: that is, all traders commence each experiment with the profit margins between 5–35%. Finally, $c_{\mathcal{R}} = 0.05$, $c_{\mathcal{A}} = 0.05$ and $,_i(0) = 0; \forall i$.

Results showing the average of 50 runs, each for 10 'days', for ZI-C and ZIP traders in markets A and B are shown in Table 1, and for markets C and D in Table 2. As predicted by Cliff and Bruten (1997), the ZI-C traders reach equilibrium in market A but fail in markets B, C, and D. However, the average transaction prices of the ZIP traders in all four markets tend toward the equilibrium price $P_0 = 200$. In markets A and B the average ZIP transaction price rapidly converges to near 200, typically within the first four trading days, and remains at that level thereafter, with low variance.

The data presented in Table 2 are less satisfactory: the initial average ZIP transaction prices are close to those of the ZI-C traders, but this is followed by a comparatively slow (yet steady) approach toward P_0, from below. To further illustrate the behavior of ZIP traders in these two markets, the experiments were continued to 30 trading days. As is clear from the $d = 30$ data, the long-term tendency of the ZIP traders is towards P_0. If the various system parameters (such as the initial distributions of profit margins, and the distributions of learning rates and momentum values) were altered, faster approach to P_0 in markets C and D could be demonstrated.

Similarly, the approach to equilibrium from below in market A is an artefact of the buyers and sellers having initial values of profit margin drawn from distributions over the same ranges of percentages, as explained by Cliff (1997). Again, the initial settings of the traders' parameters could be altered to eliminate this bias (i.e., give the sellers higher percentage profit margins than the buyers).

However, the intention here is not to demonstrate ZIP traders with optimal parameter settings: rather, the data presented here serves to demonstrate that the simple ZIP trading strategies can readily achieve results that Cliff and Bruten (1997) demonstrated to be impossible when using ZI-C traders, and are closer to those expected from human subjects or traditional rational-expectations theoretical predictions, with the same ZIP parameter values in a variety of market conditions. On these grounds at least, the minimally adaptive ZIP traders represent a significant advance on the work of G&S.

Smith's measure of allocative efficiency and G&S's measure of profit dispersion were also calculated for the ZIP traders. As with the ZI-C traders, measures of allocative efficiency for ZIP traders are typically very high (often averaging 100%). For this reason, ZIP allocative efficiency data are not very informative, and so are not shown here. However, the profit dispersion data are more revealing: time series of average profit dispersion values for both ZI-C and ZIP traders in the markets A to D are shown in Tables 3 and 4.

The profit dispersion data clearly shows that in all cases the final profit dispersion is significantly less for ZIP traders than for ZI-C traders. In markets A and B the ZIP profit dispersion falls sharply over the first four days and then levels out to a roughly constant value; in markets C and D the fall is less dramatic but could be made more rapid by appropriate alteration of the parameter-settings, as discussed previously. In Section 2.1,

d	A:ZI-C	A:ZIP	B:ZI-C	B:ZIP
1	35 ± 09	27 ± 09	27 ± 08	20 ± 06
2	33 ± 09	13 ± 09	28 ± 07	06 ± 04
4	32 ± 09	06 ± 07	27 ± 08	01 ± 01
6	33 ± 10	04 ± 05	26 ± 07	01 ± 01
10	33 ± 09	03 ± 06	27 ± 07	01 ± 01

Table 3: Profit dispersion $\mathcal{D}$ ($\times 100$) for ZI-C and ZIP traders in markets A and B. Each column shows the mean ($\pm$ standard deviation) value of $100\mathcal{D}$ for day d.

d	C:ZI-C	C:ZIP	D:ZI-C	D:ZIP
1	62 ± 17	64 ± 19	48 ± 11	38 ± 07
2	61 ± 13	58 ± 18	48 ± 10	34 ± 07
4	64 ± 15	51 ± 19	52 ± 10	27 ± 07
6	60 ± 15	47 ± 19	51 ± 08	20 ± 08
10	58 ± 12	40 ± 20	49 ± 13	10 ± 09

Table 4: Profit dispersion $\mathcal{D}$ ($\times 100$) for ZI-C and ZIP traders in markets C and D. Format as for Table 3.

it was noted that G&S state that the ZI-C profit dispersion levels are appreciably higher than those of human traders. Tables 3 and 4 demonstrate ZIP traders rapidly adapting to give profit dispersion levels that are in some cases more than a factor of ten less than those of ZI-C traders. On this basis, it seems safe to claim that the performance of the ZIP traders shown here is significantly closer to that of human traders than is the performance of ZI-C traders.

3.4 Related Work

Despite G&S's work on ZI traders having been cited approvingly in a number of texts discussing CDA markets, there appear to be very few papers that are comparable to the work described here: we know of no other critiques of G&S's work, and have found only two papers that describe artificial trading agents similar to the ZIP traders developed here. These two papers are by Easley and Ledyard (1992) and Rust, Miller, and Palmer (1992), discussed below. Cliff (1997) provides an extended critique of the market-based control literature (e.g. Clearwater (1996)), noting that the problem of incorporating bargaining mechanisms in software agents is commonly avoided by introducing centralized auctioneer processes. For this reason, no work in market-based control is reviewed here. Cliff (1997) also discusses the lack of relevant work in "artificial life" research.

Easley and Ledyard (1992) developed an analysis of hypotheses for price formation and equilibration in human markets, describing a strategy that yielded three specific predictions of human market behavior. Their trading strategy is simple, but requires more market data than do ZIP traders and can only enter into one transaction per day (ZIP traders can engage in multiple transactions: see Cliff (1997)). Easley and Ledyard's (E&L's herein) strategy is only fully effective after the first day of trading; yet it is often on the first day that the most significant shifts in behavior occur. Also, E&L's analysis relies on a simplifying assumption that is questionable in practice: they assume that, when more than one trader is interested in a transaction, the trader offering the best price is guaranteed success (1992, p.70). Several of the experimental observations E&L present contradict their theoretical predictions. Furthermore, as E&L (1992, p.87) note, their theory does not apply to experiments in which one side of the market is 'silent' (e.g., retail markets), and it doesn't predict the effects of shifts in supply and demand curves. However, swift and stable equilibration responses of ZIP traders when supply and demand alter, and the human-like equilibration of ZIP traders in 'retail' markets, are both demonstrated by Cliff (1997). Because ZIP traders give good performance in situations where E&L's work cannot be applied, it seems fair to claim that ZIP traders are both simpler than, and an advance on, the work of E&L (1992).

Rust et al. (1992) report on a series of experimental economics tournaments they organized, where other researchers were invited to submit software agents that would compete against one another in a simplified CDA market. Hence, their markets were composed of traders with *heterogeneous* strategies, and the most successful strategy was essentially parasitic: it exploited the actions of other strategies; but a *homogeneous* CDA market populated entirely by the parasite strategy exhibited manifestly suboptimal dynamics. Thus, there is no focus in Rust et al. (1992) on explicit critiques of G&S's ZI traders, or on exploring the behavior of homogeneous groups of traders in differing environments such as markets A to D.

The recent work of Epstein and Axtell (1996) has attracted much attention. This includes studies of bilateral trade between simple software agents in spatially distributed markets, but the trade mechanisms involve the exchange or bartering of two commodities: there is no money or price mechanism in their models (Epstein & Axtell, 1996, p.101), and so their work also does not bear comparison with G&S's.

3.5 Summary

The results presented in Section 3.3 demonstrated that the ZIP traders yield better results than ZI-C traders: Tables 1 and 2 showed ZI-C traders converging to equilibrium in one market but failing (as predicted by Cliff and Bruten (1997)) in another three. In contrast, the ZIP traders succeed in reaching equilibrium in all four markets. It was also demonstrated that profit dispersion is lower in ZIP trader markets than in ZI-C markets, so the ZIP results are closer to the human-trader data presented by Gode and Sunder (1993).

In addition to comparing the behavior of ZIP and ZI-C traders, we can also compare the behavior of ZIP traders to Smith's (1962) results from human subjects. In particular, Smith notes that in his excess-demand market, "...The approach to equilibrium is from below, and the convergence is relatively slow": both of these qualities are exhibited in market C by the ZIP trader results, but not the ZI-C trader results, in Table 2. Thus, the ZIP traders give results qualitatively similar to those of Smith's (1962) human subjects both in their equilibration behavior and in their modes of failure. Cliff (1997) shows ZIP traders converging to below-equilibrium prices in 'retail' markets where only sellers quote prices: Smith (1962) describes similar results with human traders. Significantly, in these markets inspired by Smith's experiments, ZIC traders either fail to give results comparable to human data, or cannot be used without revising and extending their specification.

Smith (1962) also experimented with altering supply and demand mid-way through the experiment, and with 'high-volume' markets where his human subjects were given the right to buy or sell more than one unit per day. Again, ZIP traders exhibit human-like performance in such markets, rapidly adapting to the new P_0 (Cliff, 1997).

The similarities between theoretical predictions, human data, and ZIP traders are striking and significant because of the simplicity of the trading strategies and adaptation mechanisms in the ZIP traders. While Cliff and Bruten (1997) demonstrated that G&S's ZI-C traders are *too* simple, the results in this paper indicate that the ZIP traders introduced here are simple enough to give human-like performance, but not too simple.

4 CONCLUSION

Computational trading agents can be viewed as mechanistically rigorous statements of potential models of human bargaining behaviors. G&S's work was an important contribution to the field of experimental economics, proposing an absolute lower limit on the mechanistic complexity of CDA trading agents, and demonstrating that allocative efficiency is a poor indicator of the intelligence of agents in a CDA market. However, Cliff & Bruten's (1997) critique predicts failings of ZI-C traders that were confirmed empirically. Thereby indicating a need for bargaining mechanisms more complex than simple constrained stochastic generation of prices.

The work on ZIP traders reported here should be viewed as a preliminary sketch of what forms such bargaining mechanisms might take. The ZIP traders are more complex than G&S's ZI-C 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 and profit dispersion, are clearly closer to those from human experimental markets than are the results from ZI-C traders. It is reassuring to see that very simple 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

Clearwater, S. H. (Ed.). (1996). *Market-Based Control.* World Scientific, Singapore.

Cliff, D. (1997). Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical Report HP–97–91, Hewlett Packard Research Laboratories, Bristol, UK.

Cliff, D., & Bruten, J. (1997). More than zero intelligence needed for continuous double-auction markets. Technical Report HP–97–, Hewlett Packard Research Laboratories, Bristol, UK. Submitted to CEFES98

Davis, D., & Holt, C. (1993). *Experimental Economics.* Princeton University Press.

Easley, D., & Ledyard, J. (1992). Theories of price formation and exchange in double oral auctions. In Friedman, D., & Rust, J. (Eds.), *The Double Auction Market*, pp. 63–98. Addison-Wesley.

Epstein, J. M., & Axtell, R. (1996). *Growing Artificial Societies.* MIT Press.

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

Rumelhart, D., Hinton, G., & Williams, R. (1986). Learning internal representations by error propagation. In Rumelhart, D., & McClelland, J. (Eds.), *Parallel Distributed Processing*, pp. 318–362. MIT Press.

Rust, J., Miller, J., & Palmer, R. (1992). Behavior of trading automata in a computerized double auction market. In Friedman, D., & Rust, J. (Eds.), *The Double Auction Market*, pp. 155–198. Addison-Wesley.

Smith, V. (1962). Experimental study of competitive market behavior. *J. Polit. Econ.*, *70*, 111–137.

van Montfort, G. (1997). Economic agents for controlling complex systems. Master's thesis, Delft University of Technology Faculty of Technical Mathematics and Informatics. Unpublished.

Wilson, S. W. (1994). ZCS: A zeroth level classifier system. *Evolutionary Computation*, *2*(1), 1–18.