



Zero is Not Enough: On The Lower Limit of Agent Intelligence For Continuous Double Auction Markets[†]

Dave Cliff*, Janet Bruten
Living The Vision
HP Laboratories Bristol
HPL-97-141
November, 1997

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

agent continuous
double auction,
trade

Gode and Sunder's (1993) results from using "zero-intelligence" (ZI) traders, that act randomly within a structured market, appear to imply that convergence to the theoretical equilibrium price in continuous double-auction markets is determined more by market structure than by the intelligence of the traders in that market. However, it is demonstrated here that the average transaction prices of ZI traders can vary significantly from the theoretical equilibrium value when the market supply and demand are asymmetric, and that the degree of difference from equilibrium is predictable from *a priori* probabilistic analysis. In this sense, it is shown here that Gode and Sunder's results are artefacts of their experimental regime. Following this, 'zero-intelligence-plus' (ZIP) traders are introduced: like ZI traders, these simple agents make stochastic bids. Unlike ZI traders, they employ an elementary form of machine learning. Groups of ZIP traders interacting in experimental markets similar to those used by Smith (1962) and Gode and Sunder (1993) are demonstrated, and we show that the performance of ZIP traders is significantly closer to the human data than is the performance of Gode and Sunder's ZI traders. Because ZI traders are too simple, we offer ZIP traders as a slight upward revision of the intelligence required to achieve equilibrium in continuous double auction markets.

*Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

[†]This paper describes work done while Dave Cliff held a Visiting Academic post at Hewlett-Packard Laboratories Bristol, from February to September 1996. The main results in this paper were presented at the First HP International Workshop on Interacting Software Agents, Bristol, September 1996, and at the Internet Telephony Interoperability Consortium meeting, Bristol, June 1997: we thank the participants of those meetings for their comments. Thanks also to Chris Preist for his comments on earlier versions of this manuscript.

1 Introduction

Smith (1962) demonstrated that the transaction prices of remarkably small groups of human traders, operating in experimental continuous double auction markets, rapidly approach the theoretical equilibrium price. But human beings are notoriously smart creatures: 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 Gode and Sunder (1993), whose results appear to indicate that no intelligence at all is required of the traders, so long as they are constrained not to enter into loss-making transactions.

In this paper, we use the markets introduced by Smith (1962) as inspiration for a critique of the results presented in Gode and Sunder (1993). We demonstrate that, although zero intelligence is not enough, traders with remarkably simple adaptive mechanisms can give performance very similar to that of groups of human traders.

Gode and Sunder (1993) reported results from market experiments where “zero-intelligence constrained” (ZI-C) trader-programs, that submit random bids and offers, are used to replace human traders in continuous double auction (CDA) markets. They found that the imposition of the budget constraint (that prevents ZI traders from entering into loss-making deals), is sufficient to raise the allocative efficiency of the auctions to values near 100 percent. 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.” (Gode & Sunder, 1993, p.119).

This important work has been frequently cited in the experimental economics literature. See, for example, Friedman and Rust (1992, p.xxiii), Friedman (1992, p.19), Rust, Miller, and Palmer (1992, pp.160–161, 175), Bollerslev and Domowitz (1992, pp.230–231), Cason and Friedman (1992, pp.253, 258), Kagel and Vogt (1992, pp.292, 294), Davis and Holt (1993, p.132), Roth (1995, pp.52–55, 80–81), Holt (1995, p.370), Sunder (1995, p.475), Kagel (1995, pp.570, 580), and Camerer (1995, p.674), and has even been discussed in a recent book on the philosophy of cognitive science (Clark, 1997, pp.183–184).

In this paper, we develop an analysis of the probability functions underlying CDA markets populated by Gode and Sunder’s ZI-C traders. This analysis leads to predictions of market conditions in which ZI-C traders fail to trade at equilibrium prices. We support these analytic results with results from simulation experiments: the ZI-C traders are demonstrated to fail in markets similar to those used in the early experimental economics work of Smith (1962). Thus we claim that the ZI-C traders lack sufficient rationality to exhibit human-like equilibration in CDA markets.

This leads us to develop ‘zero-intelligence-plus’ (ZIP) traders that, in the spirit of Gode and Sunder’s work, are stochastic traders with minimal intelligence. But, unlike Gode and Sunder’s programs, the ZIP traders incorporate elementary machine learning techniques to alter their behavior on the basis of experience. It is demonstrated that in experimental conditions comparable to those used by Smith (1962), where Gode and Sunder’s ZI-C traders either fail to produce human-like tendency to equilibrium or simply cannot be used, the behavior of the ZIP traders is comparable to that of human subjects. Thus, We offer ZIP traders as a revised lower limit on the degree of rationality required to trade in CDA markets.

We commence in Section 2 with brief overviews of the two papers on which our work rests. The first is Smith’s seminal 1962 paper describing his early work in experimental economics: the styles of market supply and demand functions used by Smith in that paper form the basis

of our critique of Gode and Sunder’s work. The second is Gode and Sunder’s 1993 paper itself.

Having established the background, Section 3 presents our critique of Gode and Sunder’s work and Section 4 introduces the ZIP traders and demonstrates their operation in a variety of markets. Related work is described in Section 4.5, and possible further developments of this research are discussed in Section 4.6. All the programs for our simulation experiments have been placed in the public domain (Cliff, 1997) for use by others wishing to replicate our results or experiment with variations.

2 Two landmarks in experimental economics

2.1 Smith’s early work

Smith’s 1962 paper was one of the first to discuss in detail the use of ‘laboratory methods’ in economics, conducting controlled experiments to test theoretical hypotheses or predictions. Such work has been of interest since at least the 1930’s: for historical reviews, see Davis and Holt (1993, pp.5–9) and Roth (1995, pp.4–21). Smith’s (1962) paper is reviewed here both as a means of introducing terminology and typical methods of experiment and analysis, and also to establish a more varied set of experimental scenarios than were used by Gode and Sunder: the arguments in Section 3 are supported by considering the performance of Gode and Sunder’s ZI-C traders in experimental settings similar to those used by Smith.

In a typical simple experiment, a group of human subjects are each given the means to buy some number of units of an arbitrary commodity; while another group are each given some number of units of the commodity to sell. Each buyer will be given a maximum limit price to pay for each unit of the commodity, and each seller will be given a minimum limit price below which each unit should not be sold. Typically, different buyers will be given different limit prices, as will different sellers. The manner in which limit prices are allocated determines the supply and demand for the experiment. The subjects are then allowed to buy and sell within a particular market mechanism. In the early experiments, the markets were experimentally controlled situations similar to 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. In an experimental CDA market, Active buyers may ‘shout’ (i.e., quote) bid-prices at any time, and active sellers may ‘shout’ (i.e., quote) offer-prices at any time. In the remainder of this paper, we use the word “shout” to refer to a quote from a trader, to reinforce the metaphor of open-outcry trading pits. In addition to his work on experimental CDA markets, Smith (1962) also experimented with a model ‘retail’ market where the buyers stayed silent while the sellers shouted offer-prices.

Each of Smith’s experiments were conducted over a number of trading sessions or “days” (which typically lasted around 10 minutes each). At the start of each day, the rights to sell or buy units of commodity were given to each trader. Each seller ceased to be active in the market when she had sold all the units given to her at the start of the day, and each buyer ceased to be active when she had bought as many units as she was allocated buying-rights for at the start of the day. Only active traders could shout prices. A transaction occurred whenever an active buyer decided to accept the last offer made by a seller, or an active seller decided to accept the last bid by a buyer. At the start of the next day, allocations of buying and selling rights would be redistributed to the traders, and the market would recommence trading. Usually the same traders held the same allocations throughout the sequence of trading days in any one experiment, but sometimes the market supply, or demand, or both would be altered between days by distributing different allocations of units and limit-prices.

In such experiments, hypotheses concerning the effects of market-structure on the dynamics of the market can be tested and their implications explored by varying parameters such as those affecting supply and demand, what trader actions are permissible, 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 the theoretical equilibrium price, the stability at equilibrium, the amount of potentially-available profit that is extracted from the market by the sellers, and so on. In particular, Smith (1962) monitored the *allocative efficiency* of the experimental markets, (defined as the total profit actually earned by all the traders, expressed as a percentage of the maximum possible total profit that could have been earned by all the traders), and he introduced a *price convergence* measure, α , which is calculated at the end of each trading day. For a day on which k transactions occurred at prices $p_j : j = 1 \dots k$, $\alpha = 100\sigma_0/P_0$ where P_0 is the theoretical equilibrium price given by the intersection of the supply and demand curves, and $\sigma_0 = \sqrt{\frac{1}{k} \sum_{j=1}^k (p_j - P_0)^2}$. These two metrics of market behavior will be used in the discussions below.

2.2 Gode and Sunder’s Zero Intelligence Traders

Gode and Sunder (1993) 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 CDA markets. To use currently-popular terminology, Gode and Sunder’s programs are zero-intelligence “software agents”. Gode and Sunder explored the performance of both ‘unconstrained’ and ‘constrained’ ZI traders, which they refer to using the abbreviations ZI-U and ZI-C, and they compare the results of these traders to results from human traders operating in (almost) identical experimental conditions.

As with Smith’s work, each trader is given an entitlement to buy or sell a number of units, each with a particular limit price. Bids and offers were limited to the range 1 to 200 units of arbitrary currency. For ZI-U traders, the shout-price for a unit of commodity is unconstrained across this range: there is a 1/200 probability that the value of a random shout is v for all $v \in \{1, 2, \dots, 200\}$. This means that the ZI-U traders can enter into loss-making deals: a seller of a unit of commodity could shout a price less than the given limit price, or a buyer could shout a price greater than the limit; in either case, if the shout is accepted by another trader, the agent makes a loss. In contrast, ZI-C traders are subject to a ‘budget constraint’ that prevents them from engaging in loss-making deals: a ZI-C seller can only make offers in the range between the limit price and the maximum (200); and ZI-C buyers can only make bids in the range from the minimum (1) up to their limit price. The shout prices are generated from a uniform distribution across the range for each ZI-C trader.

The experiments with both types of ZI traders were conducted using minor simplifications of the continuous double auction: traders shout prices at any time, subject to an “improvement” rule that only allows bid-prices to be higher than the current best bid, or offer-prices to be lower than the current best offer, and with a transaction cancelling any unaccepted bids and offers. The traders dealt in lot-sizes of a single unit of commodity. To accommodate the lack of intelligence of the traders, a deal was made whenever a bid and offer crossed: whenever a buyer made a bid higher than the current lowest offer, or whenever a seller made an offer lower than the current highest bid. In both cases, the transaction price is the earlier of the two shouts.

Gode and Sunder (1993) report on ZI traders with the right to buy or sell multiple units of commodity during the course of the experiment, with the lot-size for each offer, bid, and deal fixed at one unit. They state (1993, p.122) that experiments with a single unit per trader reported in (Gode & Sunder, 1992) gave similar results, although the single-unit-per-trader

studies were not concerned with the behavior of prices (Gode & Sunder, 1992, p.202). The units given to an individual seller may have different limit prices, as may the purchase rights given to an individual buyer. In one experiment, several traders were given limit prices for all their units that were extra-marginal (i.e., beyond the equilibrium point and therefore difficult to trade) while the other traders' assignments were all intra-marginal units. In the other experiments, each trader had a spread of limit prices across the range of supply or demand. To avoid requiring the traders to make a decision about which of their stock of units they should sell next, or which of their rights to buy they should exercise next, the rights to buy or sell were assigned to each agent in a pre-specified order of execution.

Differences in performance between the ZI-U and ZI-C traders, and between the ZI-C and human traders, could indicate the different extents to which overall market behavior is dependent on human intelligence or market structure:

“The difference between the performance of the human markets and that of the ZI-C markets is attributable to systematic characteristics of human traders. If ZI-C traders are considered to have zero rationality, this difference in performance would be a measure of the contribution of human rationality to market performance. On the other hand, the difference between the performance of markets that do impose a budget constraint on ZI traders and the performance of those that do not is attributable to the market discipline. Traders have no intelligence in either the ZI-U or ZI-C market: the ZI-C market prevents the traders from engaging in transactions that they cannot settle. Consequently, we can attribute the differences in market outcomes to the discipline imposed by the double auction on traders” (Gode & Sunder, 1993, p.123).

Results from five experiments are reported in the 1993 paper. Each experiment is repeated three times: once with ZI-U traders (six buyers and six sellers), once with ZI-C traders (six buyers and six sellers), and once with human subjects (six or seven buyers and six or seven sellers, all of whom were business graduate students, given an incentive to do well by having their profits included in the grading scheme for their course). Like the ZI-C traders, the humans were subject to the budget constraint: they could (or should) not enter into loss-making deals. Although the supply and demand schedules varied between experiments, in any one experiment the ZI-U and ZI-C traders had identical supply and demand schedules, and the humans schedules differed only slightly, when there was an extra seller or buyer. Nevertheless the human schedules were very similar to those of the ZI traders, and had the same equilibrium price and quantity.

Figure 1 gives an indication of the qualitative differences between the price histories of the three types of trader: in all five experiments, the same qualitative differences emerged. Prices in the ZI-U markets exhibited “...little systematic pattern and no tendency to converge toward any specific level” (Gode & Sunder, 1993, p.126). Price histories in the human markets were similar to those in Smith's early experiments: the transaction prices soon settle to stable values close to the theoretical equilibrium price, after some initial learning and adjustment. As Gode and Sunder (1993, pp.127–128) note, such price series are “...the result of subjecting profit-motivated, intelligent human traders to market discipline”, and the main question addressed in their article is: to what degree is the difference between the market activity of the human and ZI-U traders due to the intelligence of the traders, and to what degree is it due to the imposition of the budget constraint? It is this issue that the data from ZI-C traders helps resolve: they have the same budget constraint as the humans, but none of the intelligence or ability to learn from experience.

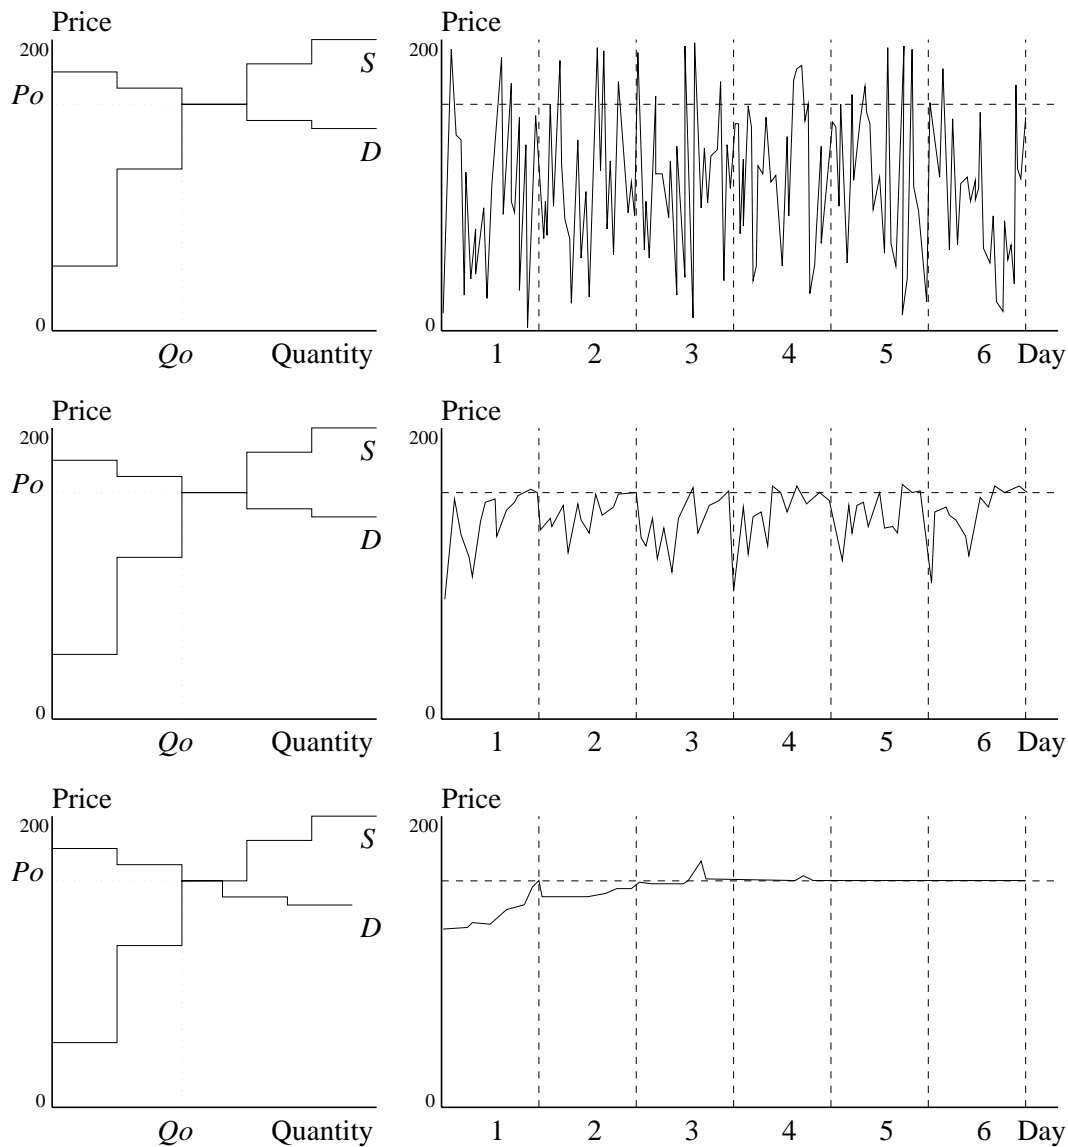


Figure 1: Typical results from one of Gode and Sunder's experiments (redrawn from Gode and Sunder (1993, Fig.4, p.127)). Top: results from ZI-U traders. Middle: results from ZI-C traders. Bottom: results from human traders. The figure to the left shows the supply and demand schedules for the experiment, the figure to the right shows the time series of transaction prices.

Gode and Sunder (1993, p.129) point out three notable features of the ZI-C transaction-price time-series. First, as would be expected from traders that have no memory or adaptation mechanisms, there is no evidence of learning from day to day. Second, the volatility of the ZI-C prices is intermediate between the highly volatile ZI-U prices and the stable human prices: the presence of a budget constraint is sufficient to shift the ZI market behavior towards that of humans. Third, despite being more volatile than the human price series, the ZI-C prices converge slowly to equilibrium during each day's trading (this is clear in the middle transaction-price time-series in Figure 1). Gode and Sunder confirm that this convergence is consistent by regression analysis of Smith's α convergence measure, averaged across the six days of a market,

versus transaction sequence number.

Gode and Sunder explain this convergence to equilibrium by reference to the progressive narrowing of the feasible range of transaction prices as more units are traded:

The left end of the market demand function represent units with higher redemption values [i.e., limit prices]. Expected values of the bids generated for these units by ZI-C traders are also higher. Therefore, these units are likely to be traded earlier than units further down the market demand function. As the higher-value units are traded, the upper end of the support of ZI-C bids shifts down. Similarly, as the lower-cost units are sold earlier in a period, the lower end of the support of ZI-C offers moves up. (Gode & Sunder, 1993, p.129).

In addition to showing human-like tendency to equilibrium, ZI-C traders also exhibit surprisingly high levels for Smith's measure of allocative efficiency. In all five experiments, the mean efficiency of human traders was over 90%, and in four of them it was over 99%. In contrast, the ZI-U traders recorded mean efficiencies of 90% in two experiments, and 77%, 49%, and 86% in the others. But the ZI-C traders scored over 99% in three experiments, and over 97% in the other two: the average efficiency for the humans was 97.9%, while for the ZI-C's it was 98.7%. Gode and Sunder (1993, pp.131–133) did not perform any statistical tests, but the difference of 0.8% between the human and ZI-C traders appears unlikely to be significant. Thus, the main message of Gode and Sunder's paper is that allocative efficiency appears to be almost entirely a product of market structure: prior to these experiments, it seemed fair to assume that the high efficiency of human markets is a consequence of human cognitive prowess; in light of Gode and Sunder's results, such assumptions are clearly highly doubtful.

Gode and Sunder close their paper with brief discussion of measurements of *profit dispersion*. This is defined (1993, p.133) as the cross-sectional root mean squared difference between the actual profits and the equilibrium profits of individual traders. The equilibrium profit of a trader is the profit the trader would realize if all units are traded at the equilibrium price. Formally, if a_i is the actual profit earned by trader i , and p_i is the theoretical equilibrium profit for that trader, then for a group of n traders the profit dispersion is given by $\sqrt{\frac{1}{n} \sum_{i=1}^n (a_i - p_i)^2}$.

Measures of profit dispersion were highest for the ZI-U traders, and lowest for the human traders. The values for ZI-C traders were closer to those of the humans than the ZI-U traders, but were generally greater in magnitude (no statistical significance tests were performed). In two of the five experiments, the human data showed a definite decline in dispersion over a number of trading days: "Without memory or learning, the ZI markets exhibit no such trend. These results suggest that, in contrast to aggregate efficiency, distributional aspects of market performance may be sensitive to human motivation and learning." (Gode & Sunder, 1993, p.134).

While Gode and Sunder's work is elegant and yields impressive results, there are some questionable aspects. For instance, the maximum and minimum values for shout-prices have to be specified in advance, which implies that *a priori* information is employed by the traders. Furthermore, it is not obvious how ZI-C traders could be used in situations such as Smith's "retail" experiment where only sellers shout (at the very least, some extension of Gode and Sunder's methods would have to be specified). However, these are relatively minor superficial criticisms. A much more fundamental critique of Gode and Sunder's work is presented in the next section.

3 Critique: Zero is not enough

Gode and Sunder’s central argument, that the structure of a double auction market is largely responsible for achieving high levels of allocative efficiency, regardless of the intelligence, motivation, or learning of the agents in the market, is not in doubt. However, in this section, the equilibrating tendencies of the ZI-C traders is questioned. Gode and Sunder state (1993, p.131):

“...the convergence of transaction price [to the theoretical equilibrium price] in ZI-C markets is a consequence of the market discipline; trader’s attempts to maximize their profits, or even their ability to remember or learn about events of the market, are not necessary for such convergence.”

This statement is demonstrated below to be incorrect. In Section 3.1 the probability distributions underlying the ZI-C markets are discussed qualitatively, noting that the average transaction prices observed in the market are determined by the intersection of the probability density functions for the buyers’ and sellers’ random bids and offers. The mean or expected value of the transaction-price distribution is shown qualitatively to be close to the equilibrium price only in situations where the magnitude of the gradient of linear supply and demand curves is roughly equal. Then, in Section 3.2, analytic results are presented that demonstrate that the expected value is equal to the equilibrium price only in certain special cases, and that the expected value differs significantly in other situations. To reinforce this result, empirical results from simulation studies are presented in Section 3.3, which are discussed further in Section 3.4.

3.1 Qualitative Discussion

Parameters relevant to the discussion that follows are shown in Figure 2.

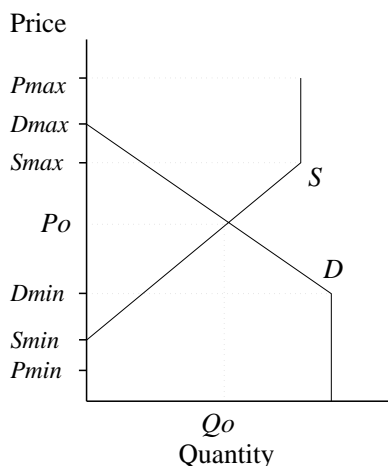


Figure 2: Parameters for qualitative discussion. Offers and bids are both subject to a maximum possible price P_{\max} and a minimum possible price P_{\min} . The Supply Curve S slopes upwards from S_{\min} to S_{\max} and the Demand Curve D slopes downward from D_{\max} to D_{\min} . The Supply and demand curves intersect at a point indicating the Equilibrium Price P_0 and the Equilibrium Quantity Q_0 . In this graph $D_{\min} > S_{\min}$ and $D_{\max} > S_{\max}$ but these conditions are not obligatory.

Because the ZI traders randomly generate bid and offer prices within given upper and lower limits, and assuming (without loss of generality) that the prices vary continuously between those

limits, it is possible to construct probability density functions (PDFs) for the prices of bids and offers shouted in the market.

In the case of ZI-U traders, the construction of the PDF is trivial: as noted by Gode and Sunder (1993, p.121), for prices in the range 1 to 200, the probability that a randomly generated bid or offer is nearest some integer value i is $1/200$ ($= 0.005$) for all $i = 1, 2, \dots, 200$. Thus a graph of the ZI-U PDF shows a uniform distribution of probabilities, as illustrated in Figure 3.




Figure 3: Probability density function (PDF) for prices ‘shouted’ by ZI-U traders: this PDF applies both to the sellers’s offer-prices and the buyers’ bid-prices.

However, the ZI-C traders have slightly more complex PDFs. Consider the case for a ZI-C seller: it can only generate an offer price between its allocated limit price, and the predetermined system-wide maximum price P_{\max} for bids or offers ($P_{\max} = 200$ in Gode and Sunder’s experiments), and the seller’s individual PDF is uniform over that range. Thus, for the market as a whole, the sellers’ supply curve acts as a lower bound on the offer prices generated at any given quantity (shown in Figure 4), and so the PDF for offer prices in the market rises from zero to some threshold value, at which it might plateau, before falling back sharply to zero at P_{\max} (Figure 5). This qualitative description can be justified as follows: there is no probability of an offer price below the lowest seller’s cost S_{\min} ; as prices higher than S_{\min} are considered, more offers are likely because there are more sellers able to bid; once prices greater than the maximum seller’s cost S_{\max} are considered, increasing prices are not increasingly likely as all sellers are able to make offers in that price range (hence the potential for plateau); and the probability of an offer above the system maximum P_{\max} is zero.

Similarly, the range of possible prices for bids from ZI-C buyers is bounded from above by the demand curve (formed by the buyers’ limit prices) and from below by the system-wide minimum price P_{\min} for bids or offers (i.e. $P_{\min} = 1$ in Gode and Sunder’s experiments): see Figure 6. Thus the PDF for bid prices in the market is zero at prices below P_{\min} , constant for prices between P_{\min} and the minimum buyer limit price D_{\min} , then falls gradually to reach zero at the highest buyer limit price D_{\max} : see Figure 7.

Having established qualitative PDFs for the *offers* generated by the sellers and the *bids* generated by the buyers, we can now consider the PDF for *transaction* prices: recall that a transaction is made when a buyer’s bid is accepted because it is higher than the current best (lowest) offer, or when a seller’s offer is accepted because it is lower than the current best (highest) bid. But the current best bid and offer must both be valid, i.e. they must come from




Figure 4: The supply curve provides a lower bound on the range of possible offer prices from ZI-C sellers at any given quantity; the system maximum P_{\max} provides an upper bound.

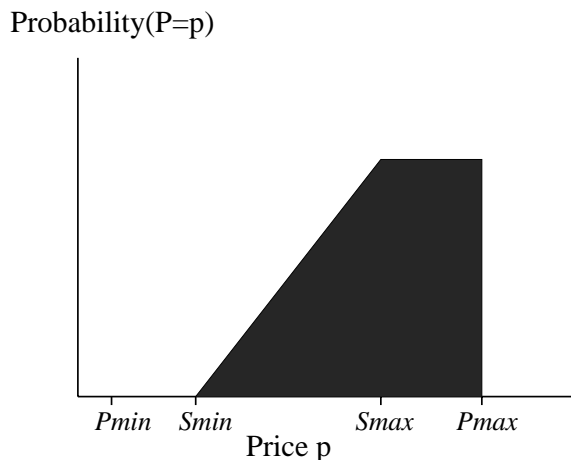


Figure 5: Qualitative probability density function (PDF) for ZI-C sellers: see text for discussion.




Figure 6: The demand curve provides an upper bound on the range of possible bid prices from ZI-C buyers at any given quantity; the system minimum P_{\min} provides a lower bound.




Figure 7: Qualitative probability density function (PDF) for ZI-C buyers: see text for discussion.

one of the PDFs in Figures 5 and 7. Thus, the PDF for transaction prices will be determined by the *intersection* of the PDFs of the offer prices and bid prices: transactions require a valid bid *and* a valid offer; despite the fact that many bids may be made at prices lower than S_{\min} , no seller can accept one; and although many offers may be made at prices higher than D_{\max} , no buyer can accept one. Figure 8 shows the intersection of the PDFs for offer and bid prices; and the corresponding PDF for transaction prices.

Finally, note that in Figure 8 the peak of the PDF, indicating the most probable transaction price(s), is determined by the intersection of the supply and demand curves; that it, the peak probability occurs at the Equilibrium Price P_0 (cf. Figure 2), and the PDF is symmetric about the peak. Thus, qualitatively at least, it would appear that near-equilibrium transaction prices

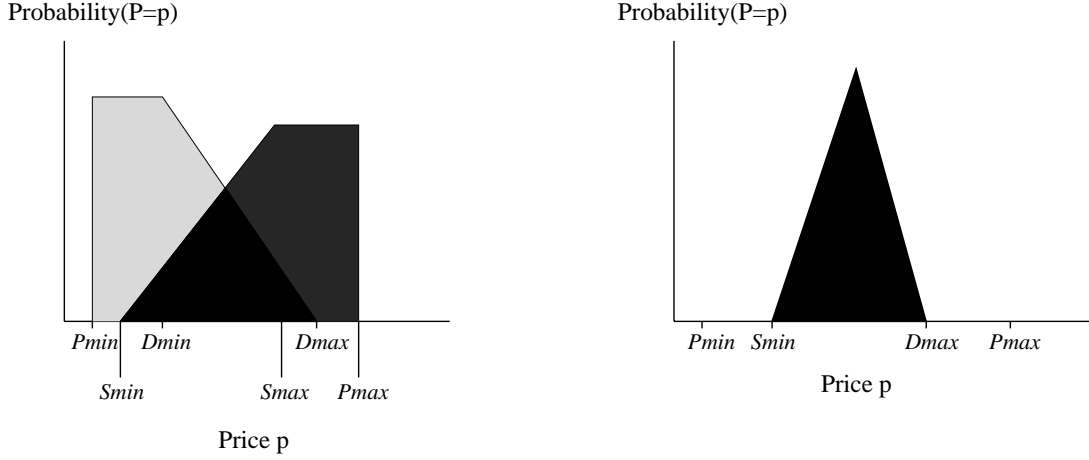


Figure 8: Left: darkest-shaded triangular area shows the intersection of the PDFs for offer and bid prices; this area indicates the PDF for transaction prices, but requires normalisation so that the area of the triangle is unity. Right: corresponding normalised PDF for transaction prices.

are expected because of the shape of the PDF for valid deals: this ZI-C system is structured *a priori* to generate mean transaction prices close to the theoretical equilibrium price.

If this is true, it is only because the transaction price PDF has a shape closely approximating an isosceles triangle, which is a consequence of the supply curve S and the demand curve D being *symmetric* about the line of constant price at P_0 (i.e., having gradients of opposite sign but approximately equal magnitude). As is proven in the next section, when the linear S and D curves have gradients opposite in sign and *identical* in magnitude, the mean transaction price is *identical* to the equilibrium price: the equilibrating tendency of the ZI-C traders is thus a consequence of the underlying probability distributions of the system. Furthermore, it is demonstrated below that, in general, the mean transaction price of ZI-C traders differs significantly from the theoretical equilibrium price.

3.2 Analytic Arguments

Let $f(p)$ denote the PDF for transaction prices. If $f(p)$ is known, then the mean or expected value $E(P)$ of the transaction prices can be calculated from the standard formula:

$$E(P) = \int_{-\infty}^{\infty} p \cdot f(p) dp \quad (1)$$

Consider the case where the supply and demand curves are symmetric (i.e., have opposite sign and equal magnitude), as illustrated in Figure 9. The corresponding PDF is shown in Figure 10. Such a market is similar to Smith's (1962) 'Chart 1', and the markets used by Gode and Sunder (1993).

The transaction-price PDF can be written as:

$$f_1(p) = \begin{cases} 0 & \text{if } p < S_{\min} \\ m_1 p + c_s & S_{\min} \leq p \leq P_0 \\ -m_1 p + c_d & P_0 \leq p \leq D_{\max} \\ 0 & p > D_{\max} \end{cases} \quad (2)$$




Figure 9: Symmetric supply and demand.




Figure 10: Transaction-price PDF for supply and demand curves of Figure 9: see Equation 2.

Let $k = D_{\max} - P_0$; note that, as the PDF is an isosceles triangle, $k \equiv P_0 - S_{\min}$ and $h_1 k = 1$. Also, from Figure 10, $m_1 = h_1/k = 1/k^2$. Therefore, substituting Equation 2 into Equation 1 gives:

$$\begin{aligned}
 E(P) &= \int_{S_{\min}}^{D_{\max}} p f_1(p) dp \\
 &= \int_{S_{\min}}^{P_0} p(m_1 p + c_s) dp + \int_{P_0}^{D_{\max}} p(-m_1 p + c_d) dp \\
 &= P_0 + \int_{-k}^0 p(m_1 p + h_1) dp + \int_0^k p(-m_1 p + h_1) dp \\
 &= P_0 + \left[\frac{1}{3} m_1 p^3 + \frac{1}{2} h_1 p^2 \right]_{-k}^0 + \left[-\frac{1}{3} m_1 p^3 + \frac{1}{2} h_1 p^2 \right]_0^k \\
 &= P_0 + \frac{1}{3} m_1 k^3 - \frac{1}{2} h_1 k^2 - \frac{1}{3} m_1 k^3 + \frac{1}{2} h_1 k^2 \\
 &= P_0 \quad \square
 \end{aligned} \tag{3}$$

Thus, Equation 3 demonstrates that when the supply and demand curves are linear, having opposite sign and equal magnitude, the mean transaction price $E(P)$ is equal to the equilibrium price P_0 .

Now consider the case for a PDF where the supply curve is flat, so that $S_{\min} = S_{\max} = P_0$ (see e.g. ‘Chart 4’ in Smith (1962)). Such a supply curve is illustrated in Figure 11, with the corresponding transaction-price PDF shown in Figure 12.

The PDF $f_2(p)$ for such a market is given by:

$$f_2(p) = \begin{cases} 0 & \text{if } p < P_0 \\ m_2 p + c_2 & P_0 \leq p \leq D_{\max} \\ 0 & p > D_{\max} \end{cases} \tag{4}$$

For $m_2 = -h_2/j$ where $j = D_{\max} - P_0$, and $c_2 = 2P_m/j^2$. Note also that because $f_2(p)$ is a PDF and a right-triangle, $h_2 j/2 = 1$, so $h_2 = 2/j$ and hence $m_2 = -2/j^2$. Substituting Equation 4 into Equation 1 gives:




Figure 11: Flat supply curve.




Figure 12: PDF for flat supply.

$$\begin{aligned}
E(P) &= \int_{P_0}^{D_{\max}} p f_2(p) dp \\
&= P_0 + \int_0^j m_2 p^2 dp + \int_0^j h_2 p dp \\
&= P_0 + \left[\frac{1}{3} m_2 p^3 \right]_0^j + \left[\frac{1}{2} h_2 p^2 \right]_0^j \\
&= P_0 + \frac{1}{3} m_2 j^3 + \frac{1}{2} h_2 j^2 \\
&= P_0 + \frac{1}{3} \cdot \frac{-2}{j^2} j^3 + \frac{1}{2} \cdot \frac{2}{j} j^2 \\
&= P_0 + \frac{j}{3} \\
&= P_0 + \frac{1}{3} (D_{\max} - P_0) \quad \square
\end{aligned} \tag{5}$$

So Equation 5 indicates that, when all the sellers have the same limit price, the expected transaction price of ZI-C traders will differ from the equilibrium price P_0 by an amount equal to one third of the difference between P_0 and the maximum buyer price, D_{\max} . So long as $P_0 \neq D_{\max}$, the expected value of the ZI-C transaction prices will differ from the equilibrium price P_0 .

Finally, consider the case of “box design” schedules (Davis & Holt, 1993, p.141), where both the supply curve and the demand curve are flat, so that $S_{\min} = S_{\max}$ and $D_{\min} = D_{\max}$. Such markets are similar to Smith’s (1962) ‘Chart 6’. If demand equals supply, then there is an indeterminate ‘price tunnel’ and the equilibrium price cannot be accurately predicted. However, if demand exceeds supply (as illustrated in Figure 13), the equilibrium price P_0 is equal to D_{\max} : the excess demand encourages price competition among buyers that will lead to bid-price increases until the maximum buyer limit price is reached. This gives a rectangular PDF, illustrated in Figure 14 and stated formally as $f_3(p)$ in Equation 6.

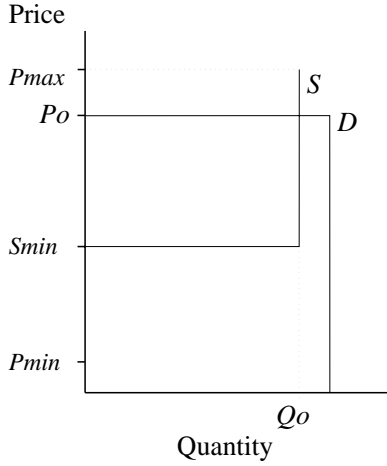


Figure 13: Flat supply and demand, with excess demand.




Figure 14: PDF for flat supply and demand.

$$f_3(p) = \begin{cases} 0 & \text{if } p < S_{\min} \\ h_3 & S_{\min} \leq p \leq P_0 \\ 0 & p > P_0 \end{cases} \quad (6)$$

Let $i = P_0 - S_{\min}$ and note that $h_3 i = 1$. Substituting Equation 6 into Equation 1 gives:

$$\begin{aligned} E(P) &= \int_{S_{\min}}^{P_0} p f_3(p) dp \\ &= S_{\min} + \int_0^i h_3 p dp \\ &= S_{\min} + \left[\frac{1}{2} h_3 p^2 \right]_0^i \\ &= S_{\min} + \frac{1}{2} h_3 i^2 \\ &= S_{\min} + \frac{1}{2} (P_0 - S_{\min}) \\ &= \frac{1}{2} (P_0 + S_{\min}) \quad \square \end{aligned} \quad (7)$$

Hence Equation 7 demonstrates that, in situations where both supply and demand are flat, and there is excess demand, then so long as $S_{\min} \neq P_0$ the expected value $E(P)$ of transaction prices will differ from P_0 .

By the same reasoning, *mutatis mutandis*, if supply and demand are flat but supply exceeds demand, then the excess supply encourages bid-price cuts and so $P_0 = S_{\min}$; the expected value $E(P)$ differs from P_0 so long as $D_{\max} \neq P_0$, and is given by Equation 8:

$$E(P) = \frac{1}{2} (P_0 + D_{\max}) \quad (8)$$

These four examples show that, for ZI-C traders, while $E(P) = P_0$ in special circumstances, in general $E(P) \neq P_0$. Similar arguments could be made for ZI-C systems with discrete rather

than continuous prices. The following section presents empirical evidence that supports the analytic argument developed here.

3.3 Simulation Studies

To test these analytic predictions, we wrote a computer simulation to study the behavior of ZI-C traders under different supply and demand schedules. The simulator was written in the C programming language: full details of the code, with sample input and output files, are given by Cliff (1997). The simulator allows for a number of ZI-C traders to be allocated roles as buyers or sellers, and to be given limit prices for their bids and offers. Results from four experiments are shown here, corresponding to the four types of supply-demand schedules examined analytically in the previous section. All experiments were run for ten trading sessions (or “days”), which continued until either eleven transactions had occurred, or no buyers or sellers were able to shout prices (because they were all unable to improve on the current best shout-prices). In each experiment, the theoretical equilibrium values are $P_0 = 200$ and $Q_0 = 6$.

Figure 15 shows a symmetric schedule where the supply and demand curves have opposite signs but equal magnitudes: there are eleven buyers each with the right to buy one unit, and eleven sellers each with a single unit to sell, The intersection of the curves indicates $P_0 = 200$. Figure 16 shows the average transaction price of each of the ten days, over 50 experiments, with lines above and below indicating the standard deviation. As is clear, the observed average transaction price is close to 200, which is the value predicted by Equation 3.




Figure 15: Symmetric supply and demand curves (gradients opposite in sign and equal in magnitude): 11 buyers and 11 sellers. Theoretical equilibrium price $P_0 = 200$; expected value of transaction prices $E(P) = 200$ from Equation 3.




Figure 16: Mean daily transaction price, averaged over 50 ZI-C experiments, for the supply and demand shown in Figure 15: the middle line is the mean value, upper and lower lines indicates the mean plus and minus one standard deviation. Horizontal axis is day-number, vertical axis is price (divided by 100). See text for discussion.

Figure 17 shows a schedule where there is a flat supply curve. Again there are 22 ZI-C traders divided equally into buyers and sellers, each with the right to trade one unit. The intersection of the curves indicates $P_0 = 200$, but Equation 5 predicts that the observed mean value of transactions will be closer to 240. Now recall that Equation 5 is for a continuous linear demand curve, while the nonlinearities in the demand curve of Figure 17 (a consequence of having only

eleven buyers) imply that the actual value of $E(P)$ may differ. By inspection, it is clear that:

$$E(P) = \left(\sum_{p=200}^{325} p \cdot g(p) \right) / \left(\sum_{p=200}^{325} g(p) \right)$$

for

$$g(p) = \begin{cases} 0 & \text{if } p < 200 \\ 5 & 200 \leq p < 225 \\ 4 & 225 \leq p < 250 \\ 3 & 250 \leq p < 275 \\ 2 & 275 \leq p < 300 \\ 1 & 300 \leq p < 325 \\ 0 & p \geq 325 \end{cases}$$

And hence the true value for the discrete nonlinear curve shown in the Figure is $E(P) = 233\frac{1}{3}$. Figure 18 then shows the average transaction price of each of the ten days, over 50 experiments, with lines above and below indicating the standard deviation. As is clear, the observed average transaction price is close to the predicted value of 233, and significantly different from the theoretical equilibrium price of 200.




Figure 17: Flat supply: 11 buyers and 11 sellers. Theoretical equilibrium price $P_0 = 200$; expected value of transaction prices $E(P) \simeq 241$ from Equation 5, which is for a continuous linear demand curve: taking into account the discrete and nonlinear nature of the curve shown here gives $E(P) = 233\frac{1}{3}$.




Figure 18: Mean daily transaction price, averaged over 50 ZI-C experiments, for the supply and demand shown in Figure 17: format as for Figure 16. See text for discussion.

Figure 19 shows a schedule where the supply and demand curves are both flat, and there is an excess of demand (eleven buyers and six sellers). The intersection of the curves indicates $P_0 = 200$, but Equation 7 predicts $E(P) = 125$. Figure 20 then shows the average transaction price of ZI-C traders for each of the ten days, again over 50 experiments, with lines above and below indicating the standard deviation. As is clear, the observed average transaction price is close to the predicted value of 125, and significantly different from the theoretical P_0 .

Finally, Figure 21 shows flat supply and demand with excess supply (six buyers and eleven sellers). Again, $P_0 = 200$ but Equation 8 predicts $E(P) = 260$. The results from z1-C traders are shown in Figure 22. Again, the observed mean transaction price is significantly closer to the predicted value of $E(P)$ than P_0 .




Figure 19: Flat supply and demand curves, with excess demand: 11 buyers and 6 sellers. Theoretical equilibrium price $P_0 = 200$; expected value of transaction prices $E(P) = 125$ from Equation 7.




Figure 20: Mean daily transaction price, averaged over 50 z1-C experiments, for the supply and demand shown in Figure 19: format as for Figure 16. See text for discussion.




Figure 21: Flat supply and demand, with excess supply: 6 buyers and 11 sellers. Theoretical equilibrium price $P_0 = 200$; expected value of transaction prices $E(P) = 260$ from Equation 8.




Figure 22: Mean daily transaction price, averaged over 50 z1-C experiments, for the supply and demand shown in Figure 21: format as for Figure 16. See text for discussion.

These four sets of simulation experiments lend strong empirical support to the analytic arguments of the previous section. In each case the empirical average transaction prices of z1-C traders are close to the value predicted from the relevant $E(P)$ equation, and in the simulations shown here the average transaction prices are only close to the theoretical equilibrium price P_0 in situations where P_0 and $E(P)$ are similar in value.

3.4 Discussion

The mathematics of Section 3.2 could be criticized for ignoring the fact that the market supply and demand curves shift after each transaction: in principle, the analysis applied only to the first transaction in each trading day. Nevertheless, there is such a good agreement between the theoretical predictions of the ZI-C traders' failure and the results from the simulations that, in practice, this criticism can be ignored.

A more subtle point is that Gode and Sunder's main claim concerned the convergence of transaction prices to equilibrium *within* a trading day: whether this happens cannot be determined from Figures 16, 18, 20, or 22, which show average transaction prices for each trading day.

To determine whether the ZI-C traders implemented here exhibit the same convergence to equilibrium as Gode and Sunder's, Figures 23 to 26 illustrate the root mean square deviation of transaction price from the equilibrium price (Smith's σ_0) calculated for each transaction sequence number. That is, for each ten-day experiment, a value $\sigma_0[1]$ is calculated from the prices of the first transaction in each of the ten days, another value $\sigma_0[2]$ is calculated from the prices of the second transaction in each day, and so on: because each day's trading with ZI-C agents is independent and identically distributed (IID), the day number is not relevant. For one experiment, this gives a vector of values $\sigma_0[i]$ where i runs from 1 to the maximum number of transactions recorded in a day. Fifty experiments give fifty such vectors, and the mean and standard deviation of the values for each element of the vector are shown in Figures 23 to 26.

As can be seen, in the symmetric market of Figure 15 and the flat-supply market of Figure 17, there is a clear reduction in deviation from equilibrium as the day progresses, indicating that the transaction prices are indeed converging on equilibrium within each trading day, as observed and explained by Gode and Sunder.




Figure 23: Horizontal axis: transaction sequence number. Vertical axis: root mean square deviation of transaction price from equilibrium price, for the symmetric market of Figure 15. Averaged over 50 experiments, each of 10 days (i.e., $n=500$). Solid line is mean; upper and lower dashed lines are at plus and minus one standard deviation respectively.




Figure 24: Horizontal axis: transaction sequence number. Vertical axis: root mean square deviation of transaction price from equilibrium price, for the flat-supply market of Figure 17. Format as for Figure 23.




Figure 25: Horizontal axis: transaction sequence number. Vertical axis: Root mean square deviation of transaction price from equilibrium price, for the excess-demand market of Figure 19. Format as for Figure 23.




Figure 26: Horizontal axis: transaction sequence number. Vertical axis: root mean square deviation of transaction price from equilibrium price, for the excess-supply market of Figure 21. Format as for Figure 23.

However, the data in Figures 25 and 26 show that the convergence to equilibrium does not occur during trading days in the ‘box design’ markets of Figures 19 and 21 respectively. There is no apparent convergence to equilibrium, and this is to be expected from consideration of the market supply and demand schedules: in these two markets, all buyers have the same limit price, and all sellers have the same limit price. Therefore each *transaction* is IID, and so there can be no correlation between transaction sequence number and transaction price. Thus, in these markets at least, there is not even a within-day convergence toward the equilibrium price.¹

3.5 Summary

The qualitative discussion of Section 3.1 led to the analytic demonstration that while there are conditions under which $E(P) = P_0$, in general the expected value of ZI-C transaction prices will differ from the equilibrium price. The empirical results presented in the Section 3.3 supported these theoretical predictions: in all the simulation studies, the theoretical equilibrium price $P_0 = 200$ and quantity $Q_0 = 6$, yet the mean daily trading price of ZI-C traders was only close to the theoretical P_0 when the supply and demand curves were symmetric: in the other cases, the mean ZI-C transaction prices deviated from the theoretical P_0 value by amounts predictable from the equations for $E(P)$. As the ZI-C traders are nothing more than stochastic systems generating random bids and offers, it would appear that the following hypothesis holds:

The mean transaction price observed in ZI-C markets can be predicted from the expected value $E(P)$ of the probability density function (PDF) given by the intersection of the sellers’ offer-price PDF and the buyers’ bid-price PDF. Only in conditions where $E(P)$ is close to the theoretical equilibrium price P_0 , will mean transaction prices

¹It is tempting to conjecture that the speed of convergence is determined by the supply and demand schedules: Figure 24 clearly shows faster convergence than figure 23, and the lack of convergence in Figures 25 and 26 could be characterized as convergence at zero speed. Possibly this is a consequence of the differences in the supply and demand schedules of the different markets: further work would be required to determine whether there is a causal link.

appear to be close to P_0 . In general, $E(P)$ and P_0 will differ, and mean transaction prices will then be at values close to $E(P)$ rather than P_0 . In brief, any similarity between ZI-C traders' transaction prices and the theoretical equilibrium price is more likely to be *coincidental* than *causal*.

Furthermore, as was demonstrated in Section 3.4, although Gode and Sunder's observation of within-day convergence of transaction prices toward the equilibrium value was replicated here in two markets (Figures 15 and 17), such convergence was not observed (and indeed is theoretically impossible) in the 'box design' markets (Figures 19 and 21).

From this it is clear that more than zero intelligence is necessary to account for convergence to equilibrium. In the next section, trading agents with slightly more than zero intelligence are introduced, and it is demonstrated that more human-like market performance is possible with remarkably little extra brain-power.

4 Zero-Intelligence-Plus (ZIP) Traders

Consider a market in which each trader is aiming for a particular level of profit: the profit margin determines the difference between the trader's limit price and shout-price. Intuitively at least, this has some appeal: initially, the only information known to a trader is the limit price(s) for the unit(s) the trader is entitled to sell or buy. In the absence of any other information, a profit-oriented buyer might shout a very low bid price, say \$0.05, even if the buyer's limit price is \$10.00: there could be a seller willing to accept the low bid, and the buyer would make a handsome profit. Similarly, a seller with a \$1.00 limit price might, if no other shouts have been made, quote a price of \$10.00 just to test the market. But, in a competitive market, with other traders holding roughly the same limit prices, these extreme shouts are unlikely to result in transactions. As soon as one shout has been made, every other trader in the market can use this to help determine a competitive price. As long as a trader can undercut a competitor and still make a profit, there is an incentive to do so. Thus, however extreme the initial shouts of greedy traders, there is a pressure on buyers to raise their bids, and on sellers to lower their offers: if the traders have set their profit margins too high, they will have to reduce them in order to remain competitive.

But it is also possible that a rational trader's profit margin will rise during a trading period. For instance, at the start of trading, a seller in possession of a unit with a \$1.00 limit price might assume that 20% profit is a reasonable level, and so intend to shout an offer of \$1.20. But, if the first few offers from competing sellers are at prices near \$10.00, and there are willing buyers at these prices, the seller would be foolish to offer at \$1.20: the seller's intended profit level could be increased forty-fold, and the resulting offer price of around \$9.00 is still likely to be accepted by a buyer.

Thus, a plausible story can be told whereby the agents in the market 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 shouts are accepted, leading to deals, or ignored. Whether such a story applies to the market behavior of humans is a matter for empirical enquiry. The notes in this section report on the development of simple mechanisms where individual traders adjust their profit margins using market price information. It is demonstrated that remarkably simple adaptive mechanisms can give performance that does not suffer from the problems affecting Gode and Sunder's ZI-C traders, discussed in the previous section. Thus, these trading agents are referred to as "zero-intelligence-plus" (ZIP) traders.

One aim of this work is identifying minimal mechanisms that endow autonomous software agents with bargaining behaviors appropriate to market-based environments. The emphasis on minimalism comes not only from a desire for computational efficiency, but also in a speculative attempt at sketching the minimum mechanistic complexity necessary and sufficient for explaining human bargaining behaviors in simple market environments.

Section 4.1 discusses at a qualitative level the conditions for raising or lowering a trader's profit margin. Section 4.2 then describes adaptive mechanisms that allow a trader's profit margin to alter over time. Results from simulation studies of ZIP traders operating in the markets used to illustrate the failure of ZI-C traders are then presented in Section 4.3. 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. Following the comparison of ZIP and ZI-C traders, Section 4.4 presents results from using ZIP traders in experimental markets similar to those used by Smith, and illustrates the dynamics of adaptation in ZIP traders. Section 4.5 describes related research, and finally Section 4.6 discusses further issues that could be explored with ZIP traders. The code and sample input and output files for the simulator system are given in Cliff (1997).

4.1 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 is whether the trader is *active* in the market (i.e., still capable of making a transaction), or *inactive* (i.e., has sold or bought its full entitlement of units, and has 'dropped out' of the market for the remainder of this trading period). The three other factors all concern the last (most recent) shout: its price, denoted by q ; whether it was a bid or an offer; and whether it was accepted or rejected (i.e., whether it resulted in a transaction or not). As discussed above, each trader maintains a profit margin, μ , which is multiplied by the limit price for a unit, λ , to determine the shout-price p . Increasing μ raises p for a seller and lowers p for a buyer. A ZIP buyer will, in principle, buy from any seller that makes an offer less than the buyer's current bid shout-price; similarly, a ZIP seller sells to any buyer making a bid greater than the seller's current offer shout-price.

For an inactive trader, there is little incentive to lower the profit margin: the trader has already successfully engaged in however many transactions it was entitled to. Even if this was a result of luck, it may as well wait for the luck to run out before it starts reducing its profit margin. But, if the trader drops out of the market and subsequently observes transactions occurring at prices which indicate that the inactive trader could have realized even higher profits, it should be able to raise its profit margin before the start of the next day. For these reasons, ZIP traders can raise their profit margins regardless of whether they are active or inactive, but only active traders reduce their margins.

When should a trader raise its profit margin? For a seller s_i , if the last shout resulted in a transaction, and s_i 's shout-price p_i was less than the transaction price q , then the indications are that s_i could have asked an even higher price and still secured a deal, so s_i should increase its profit margin μ_i . The seller's shout price could also be increased if its shout-price equals the transaction price (i.e., if $p_i = q$), because the resultant shift in the underlying supply curve should be in the seller's favor. Thus, if $p_i \leq q$, seller s_i should increase μ_i . Similarly, a buyer b_i should raise its profit margin whenever events in the market indicate that it could buy a unit for a lower price than its current shout-price p_i : that is, whenever $p_i \geq q$.

Deciding when to lower a trader's profit margin is more difficult. If sellers compete by reducing their margins whenever a buyer makes an unsuccessful bid (below the sellers' lowest

shout), the sellers would be playing into the hands of the buyers: if the buyers make a sequence of very low bids, the sellers' margins could be eroded to minimal levels while the buyers' margins remain unchanged. The reverse holds if buyers cut their profits whenever a seller makes an offer that is rejected by the buyers. Thus, if a bid is rejected, it should be the buyers that reduce their margins: in particular, any buyer who would have shouted a bid lower than the rejected value of q should reduce their margin, thereby raising their next bid. Buyers that would have shouted a bid higher than q need not reduce their margins just yet. By similar arguments, when the last shout was a rejected offer at a price q , any seller that would have shouted an offer price $p > q$ should reduce its profit margin.

But it may also be necessary for an agent to reduce its profit margin when a shout is accepted (i.e., a transaction occurs). Consider the case when a buyer makes a bid which is then accepted by a seller. If the transaction price is less than the price a seller would have shouted, then that seller should lower its profit margin because it is in danger of being undercut by the opposition (the seller that accepted the bid would have shouted a price lower than the bid). By the same reasoning, a buyer with a shout-price lower than that at which a seller's offer is accepted by some other buyer should lower its margin, so raising the price of its next bid, in order to avoid being priced out of the market.

These considerations can be summarized by the pseudo-code shown in Figure 27.

- For SELLERS:
 - if (the last shout was accepted at price q)
 - then
 1. any seller s_i for which $p_i \leq q$ should raise its profit margin
 2. if (the last shout was a bid)
 - then
 - 1. any active seller s_i for which $p_i \geq q$ should lower its margin
 - else
 1. if (the last shout was an offer)
 - then
 - 1. any active seller s_i for which $p_i \geq q$ should lower its margin
- For BUYERS:
 - if (the last shout was accepted at price q)
 - then
 1. any buyer b_i for which $p_i \geq q$ should raise its profit margin
 2. if (the last shout was an offer)
 - then
 - 1. any active buyer b_i for which $p_i \leq q$ should lower its margin
 - else
 1. if (the last shout was a bid)
 - then
 - 1. any active buyer b_i for which $p_i \leq q$ should lower its margin

Figure 27: Pseudo-code for ZIP trading strategies.

In order to test these qualitative bargaining mechanisms in (simulations of) real markets, it is necessary to specify how the profit margins of the buyers and sellers are raised or lowered.

This requires a quantitative adaptation mechanism, discussed in the next section.

4.2 Adaptation

At a given time t , an individual ZIP trader (denoted by subscript i) calculates the shout-price $p_i(t)$ for a unit j with limit price $\lambda_{i,j}$ using the trader’s real-valued profit-margin $\mu_i(t)$ according to the following equation:

$$p_i(t) = \lambda_{i,j}(1 + \mu_i(t)) \quad (9)$$

This implies that a seller’s margin is raised by increasing μ_i and lowered by decreasing μ_i , with the constraint that $\mu_i(t) \in [0, \infty); \forall t$. The situation is reversed for buyers: they raise their margin by decreasing μ_i and lower it by increasing μ_i , subject to $\mu_i(t) \in [-1, 0]; \forall t$. The aim is that the value of μ_i for each trader should alter dynamically, in response to the actions of other traders in the market, increasing or decreasing to maintain a competitive match between that trader’s shout-price and the shouts of the other traders. In order to do this, some form of adaptation or ‘update’ rule will be necessary. One of the simplest update rules in machine learning, which forms the basis of adaptation algorithms such as back-propagation in neural networks (e.g., Rumelhart, Hinton, & Williams, 1986) and reinforcement in classifier systems (e.g., Wilson, 1994, 1995), is the Widrow-Hoff “delta rule”:

$$A(t+1) = A(t) + \Delta(t) \quad (10)$$

Where $A(t)$ is the *actual output* at time t ; $A(t+1)$ is the actual output on the next time-step; and $\Delta(t)$ is the change in output, determined by the product of a *learning rate* coefficient β and the difference between $A(t)$ and the *desired output* at time t , denoted by $D(t)$:

$$\Delta(t) = \beta(D(t) - A(t)) \quad (11)$$

It is clear that, if the desired output remains constant ($D(t) = k; \forall t$), the Widrow-Hoff rule gives asymptotic convergence of $A(t)$ to $D(t)$, at a speed determined by β . The Widrow-Hoff adaptation method will be employed in the ZIP traders: when a trader is required to increase or decrease its profit margin (on the basis of the heuristics developed in Section 4.1), a ‘target price’ (denoted by $\tau_i(t)$) will be calculated for each trader, and the Widrow-Hoff rule will then be applied to take the trader’s shout-price on the next time-step ($p_i(t+1)$) closer to the target price $\tau_i(t)$. Because the shout-price is calculated using the (fixed) limit price $\lambda_{i,j}$ and (variable) profit margin μ_i , it is necessary to rearrange Equation 9 to give an update rule for the profit margin μ_i on the transition from time t to $t+1$:

$$\mu_i(t+1) = (p_i(t) + \Delta_i(t))/\lambda_{i,j} - 1 \quad (12)$$

Where $\Delta_i(t)$ is the Widrow-Hoff delta value, calculated using the individual trader’s learning rate β_i :

$$\Delta_i(t) = \beta_i(\tau_i(t) - p_i(t)) \quad (13)$$

This raises the question of how to set the target price $\tau_i(t)$. While a simple method would be to set the target price equal to the price of the last shout (i.e., $\tau_i(t) = q(t)$), this presents a significant problem. When the last shout price is very close to, or equal to, the trader’s current shout price (i.e., $p_i(t) \simeq q(t)$), the value of $\Delta_i(t)$ given by Equation 13 will be very small, or zero. Thus, traders who would have shouted prices close to $q(t)$ are likely to make negligible

alterations to their profit margins, and so will shout very similar prices when next given the opportunity. But in a competitive market, there is a need for the agents to be constantly testing the market, always pushing for higher margins. For example, if it happens that all traders are shouting prices in the range \$1.00 to \$1.05, the differences between their shouts and the transaction prices will never be more than a few cents, so they will hardly alter their shouts, and so the system will stabilize at this price range even if the true competitive equilibrium is at \$10.00. This sounds unlikely because, intuitively, it is desirable to have sellers always trying for higher prices and buyers always trying for lower prices. Thus, it is necessary for the target price to be different from the current shout or transaction price: for example, if a transaction occurs at \$1.00, a trader with a limit price of \$0.50 should aim for a price higher than \$1.00, while a buyer with a limit price of \$1.75 should aim for a target price lower than \$1.00.

There are many ways in which the target price $\tau_i(t)$ could be determined. In the current ZIP traders, the target price is generated using a stochastic function of the shout price $q(t)$, shown in Equation 14:

$$\tau_i(t) = \mathcal{R}_i(t)q(t) + \mathcal{A}_i(t) \quad (14)$$

Where \mathcal{R}_i is a randomly generated coefficient that sets the target price *relative* to the price $q(t)$ of the last shout, and $\mathcal{A}_i(t)$ is a (small) random *absolute* price alteration (or perturbation). When the intention is to increase the dealer’s shout price, $\mathcal{R}_i > 1.0$ and $\mathcal{A}_i > 0.0$; when the intention is to decrease it, $0.0 < \mathcal{R}_i < 1.0$ and $\mathcal{A}_i < 0.0$. Every time a trader’s profit margin is altered, the target price is calculated using newly-generated random values of \mathcal{R}_i and \mathcal{A}_i , which are independent and identically distributed for all traders. The use of relative increases ensures that large values of $q(t)$ are altered by greater amounts than small values of $q(t)$. For example, a shout of $q = \$10.00$ might give a seller’s target price of \$12.50 (an absolute increase of \$2.50) while a shout of \$2.00 leads to a target of \$2.50 (an absolute increase of \$0.50), but the relative increase is the same in both cases (i.e., 25%). The use of small absolute perturbations ensures that even very small shout prices lead to targets that differ by a few cents, and can be considered as random noise in the calculation of the target price.

Finally, in many applications of the Widrow-Hoff rule where the desired output $D(t)$ varies dynamically, the learning system requires ‘damping’ to prevent high-frequency oscillations around $D(t)$. Consider the case where a trader’s observations of the shouts and transactions in the market lead it to repeatedly increase its profit margin: if the next transaction to occur indicates that the profit margin is now too high, it may be premature to immediately reduce the margin; it might be better to reduce the *rate of increase* of the margin, rather than the margin itself. If the first indication that the margin should be reduced is reinforced by subsequent shouts or transactions, then eventually the rate of increase can take on a negative value (leading to reductions in the profit margin). Figuratively, the sequence of prices for shouts and transactions builds a “momentum” indicating which way the profit margin should be altered. This can easily be achieved by giving each trader a *momentum coefficient*, denoted by γ_i ($\gamma_i \in [0, 1]$) so that if $\gamma_i = 0$ the trader takes no account of past changes when determining the next change to the value of the profit margin μ_i , but with larger non-zero values of γ_i greater emphasis is accorded to past changes. Such momentum mechanisms are often employed in back-propagation neural network learning (Rumelhart et al., 1986). Equation 15 shows the general form of the equation for momentum-based updates, with $\mu_i(0) = 0; \forall i$:

$$\mu_i(t+1) = \gamma_i \mu_i(t) + (1 - \gamma_i)\Delta_i(t) \quad (15)$$

Using μ_i in place of Δ_i in Equation 12, and defining $\mu_i(0) = 0; \forall i$, gives the following update

rule, which is used in the ZIP traders:²

$$\mu_i(t+1) = (p_i(t) + \gamma_i(t))/\lambda_{i,j} - 1 \quad (16)$$

The behavior of groups of ZIP traders using the profit margin update rule of Equation 16 are illustrated in the following sections. In all the experiments reported there, \mathcal{R}_i is uniformly distributed over the range $[1.0, 1.05]$ for price increases and over $[0.95, 1.0]$ for price decreases, giving relative rises or falls of up to 5%, and \mathcal{A}_i is uniformly distributed over $[0.0, 0.05]$ for increases and $[-0.05, 0.0]$ for decreases, giving absolute alterations of up to five cents, thereby modeling a degree of uncertainty or error in the trader’s formulation of the target price. The value of each trader’s learning rate β_i is randomly generated when the trader is initialized, using values uniformly distributed over $[0.1, 0.5]$, and remains fixed for the duration of the experiment. Similarly, each trader’s momentum coefficient γ_i is randomly generated from a uniform distribution over $[0.0, 0.1]$ and remains constant for the duration of the experiment. Initial values for the μ_i profit margins of the traders are randomly generated values using a uniform distribution over the range $[0.05, 0.35]$ for sellers and $[-0.35, -0.05]$ for buyers: that is, all traders commence each experiment with the profit margins between 5 and 35 percent.

4.3 Results

To allow direct comparison, results are presented here from ZIP traders operating in the markets that were used to show the failure of ZI-C traders: the supply and demand curves for these markets were illustrated in Figures 15, 17, 19, and 21. To give a representative view of the performance of ZIP traders, all experiments are conducted with the same parameter values (listed in the previous paragraph), rather than with values optimized or ‘tuned’ to give good performance for each different market’s supply and demand curves. Results showing the average of 50 runs for ZIP traders in these four markets are shown in Figures 28, 29, 30, and 31. As can be seen by comparison with the ZI-C results in Figures 16, 18, 20, and 22, the average transaction prices of the ZIP traders in these markets are much closer to the theoretical predictions than are those of the ZI-C traders. Figures 28 and 29 clearly show average transaction prices rapidly converging to the theoretical equilibrium price of \$2.00, typically within the first four trading days and remaining at that level for the remaining days, with very little variance.

The data presented in Figures 30 and 31 are less satisfactory: the initial average transaction prices are close to those of the ZI-C traders, but this is followed by a comparatively slow (yet steady) approach to the theoretical equilibrium price, from below. To further illustrate the behavior of ZIP traders in these two markets, Figures 32 and 33 show data from experiments where there were 30 trading days, rather than 10. As is clear from these figures, the long-term tendency of the ZIP traders is towards the equilibrium price. 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 equilibrium could be demonstrated.

Similarly, the approach to equilibrium from below in Figure 28 is an artefact of the buyers and sellers having initial values of profit margin drawn from distributions over the same ranges of percentages: because the sub-marginal sellers have lower limit prices than the sub-marginal buyers, the *absolute* profit values (i.e., measured in \$) are lower for the sellers than for the buyers, and so initial transactions are more likely to occur at less-than-equilibrium prices. 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).

²Note that when there is no momentum ($\gamma_i = 0$), Equation 16 reduces to Equation 12.




Figure 28: Mean ZIP transaction prices, averaged over 50 ZIP experiments, for the supply and demand shown in Figure 15 ($P_0 = \$2.00$): format as for Figure 16. See text for discussion.




Figure 29: Mean ZIP transaction prices, averaged over 50 ZIP experiments, for the flat supply shown in Figure 17 ($P_0 = \$2.00$): format as for Figure 18. See text for discussion.

However, the intention here is not to demonstrate ZIP traders with optimal parameter-settings: rather, the data in these graphs serves to demonstrate that the simple ZIP trading strategies can readily achieve results that are 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 Gode and Sunder.




Figure 30: Mean ZIP transaction prices, averaged over 50 ZIP experiments, for the flat supply and demand, with excess demand, shown in Figure 19 ($P_0 = \$2.00$): format as for Figure 20. Trading for 10 days. See text for discussion.




Figure 31: Mean ZIP transaction prices, averaged over 50 ZIP experiments, for the flat supply and demand, with excess supply, shown in Figure 21 ($P_0 = \$2.00$): format as for Figure 22. Trading for 10 days. See text for discussion.

It is also possible to plot Smith's measure of allocative efficiency and Gode and Sunder's




Figure 32: Mean ZIP transaction prices, averaged over 50 ZIP experiments, for the flat supply and demand, with excess demand, shown in Figure 19 ($P_0 = \$2.00$): format as for Figure 20. Trading for 30 days. See text for discussion.




Figure 33: Mean ZIP transaction prices, averaged over 50 ZIP experiments, for the flat supply and demand, with excess supply, shown in Figure 21 ($P_0 = \$2.00$): format as for Figure 22. Trading for 30 days. See text for discussion.

measure of profit dispersion 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 reasons, graphs of ZIP allocative efficiency are not very informative, and so are not plotted here. However, plots of profit dispersion are more revealing: time series of average profit dispersion values for both ZI-C and ZIP traders in the four markets introduced in Figures 15, 17, 19, and 21 are shown in Figures 34 to 37 respectively.




Figure 34: Mean ZI-C and ZIP profit dispersion levels for the market of Figure 15, over 10 days. The ZI-C traders show a near-constant mean profit dispersion around 0.35, while the mean dispersion for ZIP traders falls rapidly and stabilizes at values less than 0.05.




Figure 35: Mean ZI-C and ZIP profit dispersion levels for the market of Figure 17, over 10 days. The ZI-C traders show a near-constant mean profit dispersion around 0.25, while the mean dispersion for ZIP traders falls rapidly and stabilizes at values close to 0.01.

As can be seen from the profit dispersion figures, in all cases the final (Day 10) profit




Figure 36: Mean ZI-C and ZIP profit dispersion levels for the market of Figure 19, over 10 days. The ZI-C traders show a near-constant mean profit dispersion around 0.6, while the mean dispersion for ZIP traders falls gradually from approximately 0.65 to near 0.4.




Figure 37: Mean ZI-C and ZIP profit dispersion levels for the market of Figure 21, over 10 days. The ZI-C traders show a near-constant mean profit dispersion around 0.5, while the mean dispersion for ZIP traders falls steadily from less than 0.4 to around 0.1

dispersion is significantly less for ZIP traders than for ZI-C traders. In Figures 34 and 35 the ZIP profit dispersion falls sharply over the first four days and then levels out to a roughly constant value; in Figures 36 and 37 the fall is less dramatic but could, presumably, be made more rapid by appropriate alteration of the parameter-settings, as was discussed previously. But, to reiterate, the intention here is to show the performance of ZIP traders with identical parameter settings in a variety of markets, rather than with parameter settings tuned for each market. As was discussed in Section 2.2, Gode and Sunder (1993) note that the ZI-C profit dispersion levels are lower than those of the ZI-U traders but appreciably higher than those of human traders. As is demonstrated in Figures 34 to 37, ZIP traders rapidly adapt to give profit dispersion levels that are in some cases approximately 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 in the experimental markets used here is significantly closer to that of human traders than is the performance of ZI-C traders.

4.4 Discussion

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. The symmetric market in Figure 15 is clearly comparable to the markets used in some of Smith’s early experiments (e.g., Charts 1 to 3 in Smith (1962)). The flat-supply market in Figure 17 is comparable to Smith’s (1962) Chart 4), and the excess-demand market in Figure 19 is comparable to Smith’s (1962) Chart 6). 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 by the ZIP trader results in Figures 30 and 32 but not the ZI-C trader results in Figure 20.

In the ZIP experiments shown so far, the supply and demand schedules have remained fixed for the duration of the experiment. However Smith (1962) also experimented with dynamic changes in supply or demand: in some of his experiments, at the end of a trading day a new set of limit prices was distributed to the buyers, sellers, or both. Typically, the human traders would

adapt, converging to the new market equilibrium values. This rapid, robust, and decentralized adaptation is one of the attractions of using the continuous double auction as a market organization. Thus, it is important to explore the behavior of ZIP traders when supply or demand alter (either increase or decrease).

Figure 38 shows a transaction-price time-series from one experiment which uses the symmetric market of Figure 15 for the first ten days. At the end of Day 10, the demand curve is shifted upwards by adding \$0.50 to each buyer’s limit price (P_0 increases to \$2.25), and the experiment continues for another five days. Figure 39 shows the average results from 50 such experiments. Similarly, Figure 40 shows transaction prices from one experiment where the symmetric market of Figure 15 is again used for the first ten days, but an increase in supply is then imposed by subtracting \$0.50 from each seller’s limit price (P_0 decreases to \$1.75) and trading continues for another five days. Figure 41 shows the average results from 50 such experiments. These figures clearly demonstrate that groups of ZIP traders are capable of rapidly adjusting to new equilibrium values resulting from changes in supply or demand.




Figure 38: Transaction-price time series for one experiment with a sudden increase in demand. Initial market is illustrated in Figure 15 ($P_0 = \$2.00$). After 10 trading days, demand is increased ($P_0 = \$2.25$) and the experiment continues for another 5 days. See text for discussion.




Figure 39: Mean ZIP transaction prices, averaged over 50 increased-demand experiments.

4.5 Related Work

As was noted earlier, Gode and Sunder’s work on ZI traders has been cited approvingly in a number of texts discussing continuous double auction markets. Despite this, there appear to be very few papers that are comparable to the work described here: we know of no other critiques of Gode and Sunder’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 et al. (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 the field known as “artificial life”.




Figure 40: Transaction-price time series for one experiment with a sudden increase in supply. Initial market is illustrated in Figure 15 ($P_0 = \$2.00$). After 10 trading days, supply is increased ($P_0 = \$1.75$) and the experiment continues for another 5 days. See text for discussion.




Figure 41: Mean ZIP transaction prices, averaged over 50 increased-supply experiments.

Easley and Ledyard (1992) consider several theories for price formation and equilibration, attempting to explain how human traders converge to equilibrium. They introduce a mathematical notation which they use to describe specific hypotheses concerning trading strategies and equilibration in double-auctions; a number of analytic proofs then lead to three specific predictions, which they test by comparison to data from human experiments. Their trading strategies are simple mechanisms which rely on a memory of data from past trading days. Specifically, Easley and Ledyard’s trading strategies use the following information: the lowest-priced offer or transaction in the previous day’s trading; the highest-priced bid or transaction in the previous day’s trading; the most recent bid, offer and transaction prices in the current day’s trading; the time remaining to the end of the current trading day; and an indicator of whether the agent has traded in the current day (their traders are designed to trade only one unit per day, so this indicator is similar to the way in which the ZIP traders cease to be active once they have traded all their entitlement; however, ZIP traders may enter into more than one transaction per day). Clearly, Easley and Ledyard’s trading strategies could require more memory than a ZIP trader, and also their 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. Easley and Ledyard’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 buyer with the highest shout price or the seller with the lowest shout price is guaranteed the deal (Easley & Ledyard, 1992, p.70). Despite (or possibly because of) this, several of the experimental observations they present contradict their theoretical predictions. Furthermore, as Easley and Ledyard (1992, p.87) note, their theory does not apply to experiments in which one side of the market is not allowed to bid or offer (e.g., retail markets), and it doesn’t predict the effects of shifts in supply and demand curves. In Section 4.4, we showed the responses of ZIP traders when supply and demand alter, and the human-like equilibration of ZIP traders in ‘retail’ markets is discussed in detail by Cliff (1997). Because ZIP traders give good performance in situations where Easley and Ledyard’s work cannot be applied, it seems fair to claim that ZIP traders are both simpler than, and an advance

on, the work of Easley and Ledyard (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 double auction. The double auction was simplified by synchronizing it into a two-stage process that was iterated several times per trading day. In the first stage, all traders simultaneously shout a bid or offer, and these shouts are distributed to all traders. In the second stage, the trader with the highest current bid and the trader with the current lowest offer are given the option of entering into a transaction: they can either agree a deal, or refuse. In addition to the array of bids and offers, each trader has access to public information which includes: the number of buyers; the number of sellers; the identities of the traders; the number of rounds (experiments), periods (days per experiment) and timesteps (iterations per day); the number of units each agent will have; and the distribution from which the unit limit prices are generated (Rust et al., 1992, p.164). The trading agents were allowed to be both buyers and sellers, although some researchers submitted seller-only or buyer-only strategies. A number of tournaments were held, and the different strategies were ranked in order of the profits they generated. Few details are given of the specifications of the different strategies, so a detailed comparison with ZIP traders is difficult. However, a key difference between these tournaments and the ZIP (and ZI-C) experiments is that the tournaments involved *heterogeneous* groups of traders. Traders with radically different strategies could compete in the same market, and much of the focus in Rust et al. (1992) is on the way in which the different trading strategies interacted, both with a fixed number of different strategies and in ‘evolutionary’ tournaments where the relative numbers of the different trading strategies altered over time, so more profitable strategies became more numerous than less profitable ones. The ‘population dynamics’ of the tournaments occupy much of the discussion:

“We find that the top-ranked programs yield a fairly “realistic” working model of a [double auction] market in the sense that their collective behavior is consistent with the key “stylized facts” of human experiments. We also find that a very simple strategy is a highly effective and robust performer in these markets. This strategy was able to outperform more complex algorithms that use statistically based predictions of future transaction prices, explicit optimizing principles, or sophisticated “learning algorithms”. The basic idea behind the approach can be described quite simply: *wait in the background and let others do the negotiating, but when bid and [offer] get sufficiently close, jump in and “steal the deal”*. However, the results of our evolutionary tournaments show that when too many other traders try to imitate this strategy, market efficiency can fall precipitously. . . . Specifically, if too many traders “wait in the background”, little information is generated until just before the end of the trading period. This tends to produce “closing panics” as traders rush to unload their [units] in the final seconds of the trading period, resulting in failure to execute all potentially profitable transactions.” (Rust et al., 1992, p.157, original emphasis).

Thus, there is no focus in Rust et al. (1992) on explicit critiques of Gode and Sunder’s ZI traders, or on exploring the behavior of homogeneous groups of traders in particular market environments such as the symmetric, flat supply, excess-supply ‘box’, excess-demand ‘box’, increased-demand symmetric, or increased-supply symmetric used with ZIP traders in Sections 4.3 and 4.4, or the ZIP ‘retail’ markets described in Cliff (1997). Furthermore, the reproductive success of the “wait in the background” trader strategy indicates that the ‘evolutionary’ tournaments can favor trading strategies that, when used to form homogeneous groups of traders, can give rise to market dynamics that are manifestly sub-optimal.

The recent work of Epstein and Axtell (1996) 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 Gode and Sunder's.

4.6 Further Work

While the results presented in the previous section indicate that the ZIP bargaining mechanisms give results more comparable to human traders than do Gode and Sunder's ZI-C traders, there are many possible ways in which this work could be extended.

First, the rationale for the ZIP mechanisms comes from the qualitative arguments of Section 4.1, and while the results are promising, it would be more satisfactory to develop a more rigorous, analytic treatment of these mechanisms. The qualitative rationale could perhaps be supported by game-theory analysis, or the algorithmic complexity could be analysed both in time and in space (e.g. costs for storage and network bandwidth). Furthermore, it would be attractive to develop proofs concerning the convergence to equilibrium of ZIP systems.

Also, the demonstrations of the ZIP traders come from simulations of minimally simple markets, similar to those used in Smith's early experiments. There are a variety of ways in which the complexity of the market environments could be increased, which may reveal the need for revisions or extensions to the basic ZIP mechanisms introduced here. A recent masters thesis by van Montfort (1997) describes experiments using our ZIP traders in spatially structured environments (i.e., segmented markets), and studying ZIP traders in markets where intertemporal arbitrage is permitted.

The current ZIP traders can deal with the right to sell or buy one or more units of commodity with the same limit price. Relatively straightforward extensions include endowing the agents with the right to buy or sell multiple units of commodity, where each agent's units have different limit prices, possibly also with the lot-size of each deal being chosen by the agents. A natural next step would then be to have multiple types of commodity, with the possibility of substitution between different commodity types.

A significant issue to examine is the extent to which the market dynamics are affected by the division of time into discrete 'days': in many real markets this assumption may be untenable, and it may be the case that agents 'drop out' of the market and re-enter it in an asynchronous fashion, for varying periods of time. The introduction of delays and noise into distributed markets is also likely to have a significant impact on their dynamics. In particular, delays and noise introduce uncertainty and risk: received signals might be incorrect, through corruption by noise or as a consequence of being out-of-date. Because of this, it may be necessary for the traders to reason about their 'beliefs' concerning the reliability of the data used in making trading decisions.

As a complement to the current system where ZIP traders buy and sell in a commodity market, ZIP traders could be developed for use in asset markets, where the items bought and sold can generate an income (the dividend stream) while they are owned (see, e.g., Davis and Holt (1993, p.162ff.) and Sunder (1995)). Moreover, the current ZIP market is a spot market: it would also be interesting to evaluate the performance of ZIP traders in derivative markets.

It seems very probable that, as market organizations with fewer simplifying constraints are used and as the range of possible actions available to the traders increases, more complex decision and adaptation mechanisms will be required. An obvious first approach would be to introduce "higher-order" adaptation mechanisms so that values which are currently parameters

for each agent become variables. That is, the values for the learning rate (β in Equation 13) and momentum (γ in Equation 15) for each agent could be varied dynamically on the basis of that agent's experiences in the market. Also, other variables may be introduced into the adaptation and bargaining mechanisms: there are a number of variables that the current ZIP traders do not take account of which a human trader might use to determine more profitable prices. Examples include: whether there are more buyers than sellers (or more offers than bids shouted) and vice versa; the time remaining until the end of the trading period; predictions of cyclical fluctuations in supply and demand; the average prices of the competition (to allow aggressive or predatory pricing, under-selling or over-bidding to attack the competition); and so on. Also, the current ZIP traders are specified as discrete-time processes, but in more realistic (i.e., more complex) markets, it is likely that continuous-time processes will be required.

Although more traditional machine learning techniques may also be usefully employed, recent work in biologically-inspired computing (so-called "artificial life") has seen the development of a number of adaptation mechanisms that could be employed in automatically adapting or tuning trading agents. Such techniques include the wide variety of neural-network learning algorithms (see e.g. Rumelhart and McClelland (1986) and Hertz, Krogh, and Palmer (1991)), and evolutionary approaches such as genetic programming and classifier systems (see e.g. (Davis, 1991)). This may allow further exploration or strengthening of the links between evolutionary and economic dynamics (see e.g. Hodgson (1993) and Vromen (1995)). In all cases, the profit accrued by an agent could be used as an obvious reinforcement payoff of "fitness" or "reward".

4.7 Summary

This section commenced with the simple qualitative arguments, presented in Section 4.1, for how a minimally intelligent trader might operate, and Section 4.2 then discussed some elementary quantitative adaptation mechanisms taken from the machine learning literature. Together, these strategy and adaptation mechanisms define the current ZIP traders.

The results presented in Section 4.3 demonstrated that the ZIP traders yield better results than ZI-C traders: Section 3.3 showed ZI-C traders converging to equilibrium in one market but failing (as predicted) in another three: the ZIP traders do not fail to reach equilibrium in any of these 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).

The ZIP traders were demonstrated to give results qualitatively similar to those of Smith's (1962) human subjects: even the modes of failure are similar in ZIP and human traders. That is, like humans, ZIP traders showed a slow approach to equilibrium from below in the excess-demand markets of Figures 30 and 32. Also, Cliff (1997) shows ZIP traders exhibiting a convergence to below-equilibrium prices in the 'retail' markets where only sellers shout prices: Smith (1962) describes similar results with human traders. Significantly, in these markets inspired by Smith's experiments, Gode and Sunder's ZI-C traders either fail to give results comparable to human data, or cannot even 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.

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 Section 3 demonstrated that ZI-C traders are *too* simple, the results

in this section indicate that ZIP traders are simple enough to give human-like performance, but not too simple. Having established these baseline results, Section 4.6 sketched out possibilities for extending this work. Clearly, there is much further work that could be done.

5 Conclusion

Computational trading agents can be viewed 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 (1993) or Easley and Ledyard (1992), for conveniently testing theories concerning the behavior of humans in different market structures and conditions.

Gode and Sunder's work was an important contribution to the field of experimental economics, providing an absolute lower limit on the mechanistic complexity of trading agents, and demonstrating that allocative efficiency is a poor indicator of the intelligence of agents in a double-auction market. However, the critique presented in Section 3 indicates that some of the tendencies of ZI-C traders towards theoretical equilibrium values are predictable from *a priori* analysis of the probability functions of the system. There is a sense in which the ZI-C simulation experiments (both our own, and Gode and Sunder's) are superfluous: the mathematical analysis predicts both the failures and the (apparent) successes of markets populated by ZI-C traders. The failings of the ZI-C traders indicates a need for bargaining mechanisms more complex than the simple stochastic generation of bid and offer prices.

The work on ZIP traders, reported in Section 4, 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-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

- Arthur, W. B. (1993). On designing economic agents that behave like human agents. *Evolutionary Economics*, 3, 1–22.
- Bollerslev, T., & Domowitz, I. (1992). Some effects of restricting the electronic order book in an automated trade execution system. In Friedman, D., & Rust, J. (Eds.), *The Double Auction Market: Institutions, Theories, and Evidence*, pp. 221–252. Addison-Wesley, New York.
- Camerer, C. (1995). Individual decision making. In Kagel, J. H., & Roth, A. E. (Eds.), *The Handbook of Experimental Economics*, pp. 587–703. Princeton University Press, Princeton, NJ.
- Cason, T., & Friedman, D. (1992). An empirical analysis of price formation in double auction markets. In Friedman, D., & Rust, J. (Eds.), *The Double Auction Market: Institutions, Theories, and Evidence*, pp. 253–284. Addison-Wesley, New York.
- Clark, A. J. (1997). *Being There: Putting Brain, Body, and World Together Again*. MIT Press Bradford Books, Cambridge MA.

- Clearwater, S. H. (Ed.). (1996). *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, Singapore.
- Cliff, D. (1997). Minimal-intelligence agents for bargaining behaviors in market-based environments. HP Labs technical report HP-97-91, Hewlett Packard Research Laboratories, Bristol, England.
- Davis, D. D., & Holt, C. A. (1993). *Experimental Economics*. Princeton University Press, Princeton, NJ.
- Davis, L. (Ed.). (1991). *Handbook of Genetic Algorithms*. Van Nostrand Reinhold, New York.
- 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: Institutions, Theories, and Evidence*, pp. 63–98. Addison-Wesley, New York.
- Epstein, J. M., & Axtell, R. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. MIT Press Bradford Books, Cambridge MA.
- Friedman, D. (1992). The double auction market institution: A survey. In Friedman, D., & Rust, J. (Eds.), *The Double Auction Market: Institutions, Theories, and Evidence*, pp. 3–26. Addison-Wesley, New York.
- Friedman, D., & Rust, J. (Eds.). (1992). *The Double Auction Market: Institutions, Theories, and Evidence*. Addison-Wesley, New York.
- Gode, D. K., & Sunder, S. (1992). Lower bounds for efficiency of surplus extraction in double auctions. In Friedman, D., & Rust, J. (Eds.), *The Double Auction Market: Institutions, Theories, and Evidence*, pp. 199–219. Addison-Wesley, New York.
- Gode, D. K., & Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1), 119–137.
- Hertz, J., Krogh, A., & Palmer, R. G. (1991). *An Introduction to the Theory of Neural Computation*. Lecture Notes Volume I. Addison Wesley.
- Hodgson, G. M. (1993). *Economics and Evolution: Bringing life back into economics*. Polity Press, Cambridge.
- Holt, C. A. (1995). Industrial organization: A survey of laboratory research. In Kagel, J. H., & Roth, A. E. (Eds.), *The Handbook of Experimental Economics*, pp. 349–443. Princeton University Press, Princeton, NJ.
- Kagel, J. H. (1995). Auctions: A survey of experimental research. In Kagel, J. H., & Roth, A. E. (Eds.), *The Handbook of Experimental Economics*, pp. 501–585. Princeton University Press, Princeton, NJ.
- Kagel, J. H., & Vogt, W. (1992). Buyer’s bid double auctions: Preliminary experimental results. In Friedman, D., & Rust, J. (Eds.), *The Double Auction Market: Institutions, Theories, and Evidence*, pp. 285–306. Addison-Wesley, New York.
- Roth, A. E. (1995). Introduction to experimental economics. In Kagel, J. H., & Roth, A. E. (Eds.), *The Handbook of Experimental Economics*, pp. 3–109. Princeton University Press, Princeton, NJ.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In Rumelhart, D. E., & McClelland, J. L. (Eds.), *Parallel Distributed Processing, Volume 1: Foundations*, pp. 318–362. MIT Press Bradford Books, Cambridge MA.
- Rumelhart, D. E., & McClelland, J. L. (Eds.). (1986). *Parallel Distributed Processing, Volume 1: Foundations*. MIT Press Bradford Books, Cambridge MA.
- 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: Institutions, Theories, and Evidence*, pp. 155–198. Addison-Wesley, New York.
- Smith, V. L. (1962). An experimental study of competitive market behavior. *Journal of Political Economy*, 70, 111–137.

- Sunder, S. (1995). Experimental asset markets: A survey. In Kagel, J. H., & Roth, A. E. (Eds.), *The Handbook of Experimental Economics*, pp. 445–500. Princeton University Press, Princeton, NJ.
- van Montfort, G. (1997). Economic agents for controlling complex systems. Master's thesis, Delft University of Technology Faculty of Technical Mathematics and Informatics. Unpublished.
- Vromen, J. J. (1995). *Economic Evolution*. Routledge, London.
- Wilson, S. W. (1994). ZCS: A zeroth level classifier system. *Evolutionary Computation*, 2(1), 1–18.
- Wilson, S. W. (1995). Classifier fitness based on accuracy. *Evolutionary Computation*, 3(2), 149–175.