

Impact of Activation in a Zero-Intelligence Trader Double Auction Market

Kenneth Comer

Volgenau School of Engineering
George Mason University
Fairfax, VA
(703) 675-7053
kcomer@gmu.edu

Andrew Loerch

Volgenau School of Engineering
George Mason University
Fairfax, VA
aloerch@gmu.edu

Robert Axtell

George Mason University
Fairfax, VA
Santa Fe Institute
Santa Fe, NM
rax222@gmu.edu

ABSTRACT

A model of a double auction market of zero-intelligence traders was replicated as an agent-based model using the same market supply and demand curves. The original results were reproduced, and these results and other behavior of the model were examined under different schemes of agent activation, both exogenous and endogenous. While the qualitative differences were typically minor, there were statistically significant differences in all the measures of all the markets in the original research and important divergence in the extended evolution of the simulation. These differences have important implications for all follow-on replications of a zero-intelligence trading model.

Categories and Subject Descriptors

F.3.3 [Logics and Meanings of Programs]: Studies of Program Constructs – *program and recursion schemes*.

General Terms

Algorithms, Performance, Design, Experimentation, Standardization, Theory, Verification.

Keywords

agent-based simulation, activation, updating, model replication, standardization, market design, zero-intelligence traders.

1. BACKGROUND

Finance is an area of high activity for complexity science and agent-based models. It was one of the primary motivations behind the founding of the Santa Fe Institute [4]. Agent-based models, with their many independent decision-makers, are excellent surrogates for traders in a securities market. Agents can be infused with a number of different strategies, and global information can be made available either market-wide or differentially to only select traders.

One of the simplest market models is the "zero-intelligence trader" or ZIT model. Pairs of traders are chosen from a larger body. In the most straightforward ZIT models, traders trade a single commodity. They cannot access market-wide parameters such as the last trade price or the trade price history or even the details of their counterparty's financial position. The traders are not completely devoid of knowledge: the sellers know their own cost of acquisition, and the buyers know what future price at which they can expect to liquidate the asset. (The latter might seem a bit artificial, but is analogous to the book value of assets or the

surrender value of a bond.) The simplicity of the ZIT model invites excursions on model format and design, such as studying the impact of activation to see if different activation (or updating) schemes result in different outcomes or even different policy recommendations.

2.1 On the Choice of the Gode and Sunder Model

The most referenced ZIT model was introduced by Gode and Sunder [2] in an article entitled "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality." Nearly 1200 scholarly articles have referenced Gode and Sunder over the past two decades. They were initially researching whether a rule-based double auction market simulation would show the same market success as an experimental market of actual individuals. They used graduate students incentivized by academic grade credits to simulate profit-motivated traders. They then simulated two double auction markets to compare with the real-world experiment.

2. BOUNDED ZIT MODEL DESCRIPTION

Both simulations began with a small number of traders: six buyers and six sellers. Traders trade one 'share' at a time. One simulation was unbounded, with the buyers and sellers making offers randomly selected between 0 and 200. The more rational simulation was termed a 'bounded' or constrained simulation. The buyers have a 'supply' curve in which the cost for their next share to be sold is determined by an escalating price curve. The sellers likewise have a redemption price, at which they may liquidate any item they buy. This redemption price curve decreases depending upon how many shares the buyers have already. After each trade, buyers and sellers calculate their profit. Buyers subtract the cost from the trade price, and sellers subtract the trade price from the redemption price. Buyers and sellers are bounded in that they are not allowed to make an offer that would lose money.

Gode and Sunder made a three simplifications to a double auction model:

- Only one unit was traded at a time.
- Once a trade took place, all outstanding offers were canceled.
- If bid and ask offers crossed (seller asked more than the buyer bid or vice versa), the price was set by that of the earliest offer.

Zero Intelligence Traders, 8 Runs - Activation Type: random
Bounded Trading, Double Auction Market

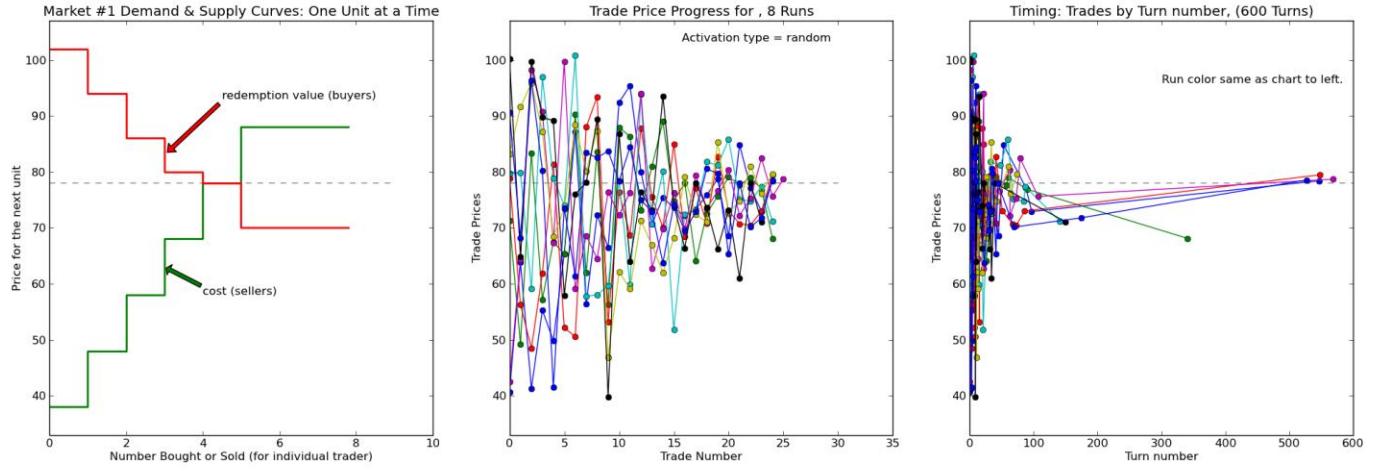


Figure 1. Market 1 Trade Price vs. Trade Number and vs. Turn

Buyers are informed ‘privately’ of the redemption value of each share. This value, v_p , depends on the number of shares the individual buyer has already bought. The buyer knows his own demand curve, but the market demand curve is not available to any trader. Similarly, sellers are endowed with a supply curve that represents the cost, c_i , of the i^{th} unit sold. This supply curve applies to each individual seller and the market supply curve is also not known to any trader. Each trade, therefore, created a profit. For the seller the profit is the net of the price and the cost, $p - c_i$. Similarly, the buyer’s profit is the net of the redemption value and the price, $v_i - p$.

Buyers and sellers form offers at a rate and in a sequence determined by the activation scheme. All buyers have the same individual demand curve, and all sellers have the same individual supply curve. The offer for buyers is a random value between 0 and their current redemption value, v_i . The offer for sellers is a random value between their cost, c_i , and 200. This is what was meant by the bounded market. The unbounded market was also examined, but that is not considered here. (Nor is the experiment using graduate students.)

Gode and Sunder conducted six runs of the bounded market, with all values reset at the beginning of each run. The runs were terminated after 30 seconds. Gode and Sunder examined five markets, or five sets of supply and demand curves. These curves were described in market-by-market graphs beside the trade price series. For the first four markets it was possible to estimate these values by inspection, but the fifth market had supply and demand curves with a structure with too fine a grain to reliably estimate. Only markets one through four were replicated here.

3. MODEL REPLICATION

Working in Python, we were able to create a double-auction model in which the traders behave in the manner described in the source article. In order to perform diagnostics, it was necessary to impose some metrics on the dynamic processes of the model. We introduced the concept of a *turn*, which we define lasting as long as one full population of traders have generated offers. A turn, therefore, is driven by events and not by time. This deviates somewhat from the source article, but allows side-by-side comparison of a variety of activation schemes (see below).

Once the turn in which trades take place is recorded, a price series

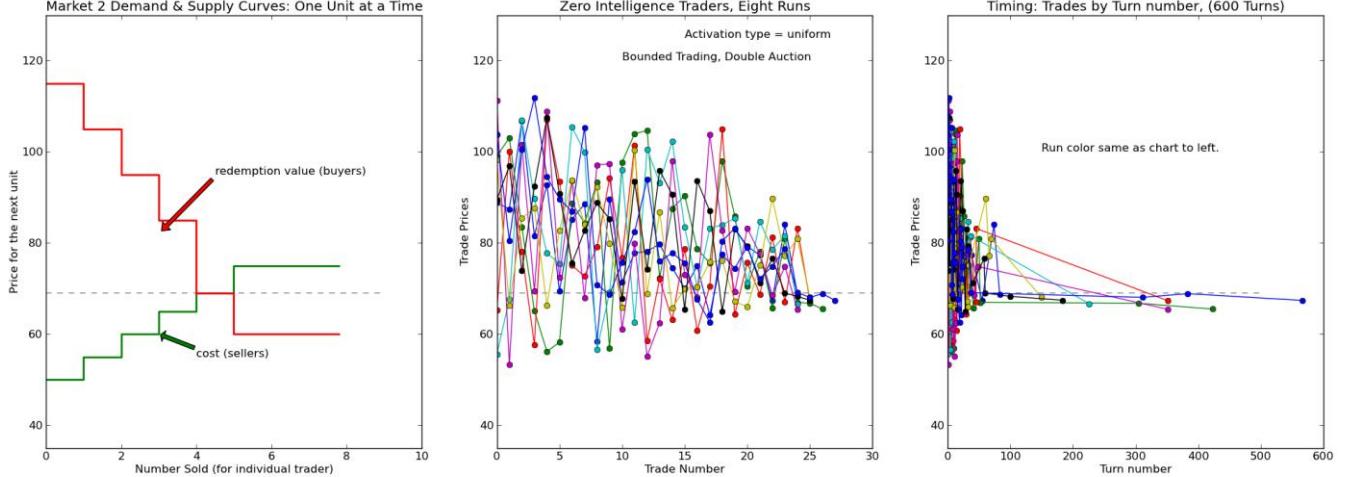


Figure 2. Market 2 Trade Price vs. Trade and vs. Turn

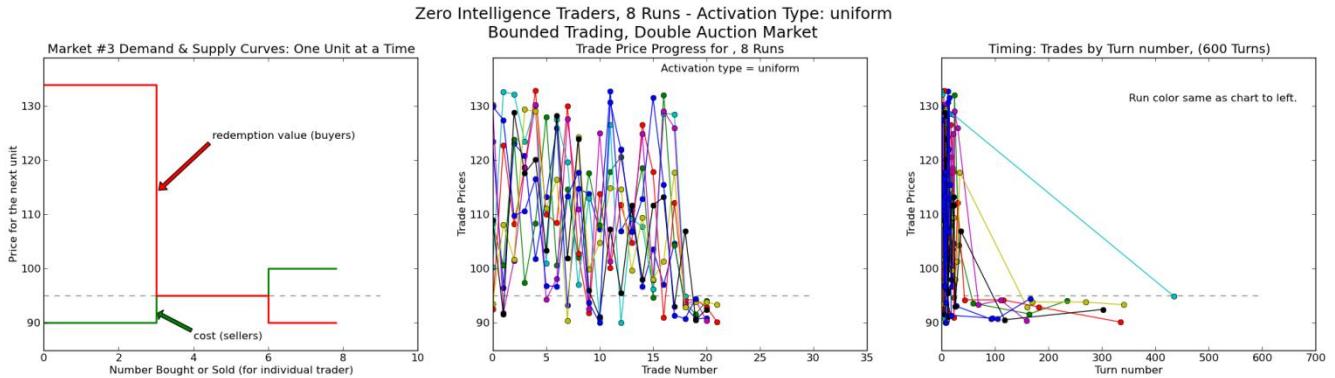


Figure 3. Market 3 (Uniform) Trade Price vs. Trade and vs. Turn

of trades can be observed in market time instead of trade time. The Gode and Sunder paper plotted trade price per trade number. Thus, they did not observe the fact that later trades occurred much later in a run, after many, many offers had been made. See Figure 1 for a depiction of this dynamic behavior for Market 1.

Figure 1 also shows a number of other aspects of our market model. Instead of stopping after 30 seconds of execution, we have chosen to stop after a constant number of turns. For this graphic, we chose 600 turns, but in the full experiments we ran the market out to 5000. Even with these extended runs there are still trades taking place. That is, even after many turns and many offers are generated there is still one buyer or seller who has redemption or cost set just above or below the market-clearing price.

Figure 2 shows a market with the asymmetry in the opposite direction – a steeper demand curve and a shallower supply curve. In both cases the trades approach the market clearing price from the direction of the steepest curve. In Market 1, they approach from below because the supply curve is steeper. In Markets 2 and 3, trades arrive at the market clearing price from above because the demand curve is steeper.

Gode and Sunder were investigating how much of the rationality associated with human traders could be attributed to human decision-making motivated by profit and intelligence and how much is due to simple market discipline – the requirement that a seller can't sell below cost and a buyer can't buy above redemption value. While the bounded market's appears to be in between the random and the human market (by inspection), and the bounded market appears to converge to the same equilibrium price as the human market (determined by a regression of the bounded market curves, averaged over five runs), Gode and Sunder measured the outcome with two quantitative measures: market efficiency and wealth distribution.

In the market evolution figures the supply and demand curves for each market was determined from the reference paper, but the price time series results were from our own replication of this double-auction model coded in Python.

4. ALTERNATIVE ACTIVATION SCHEMES

In replicating this model, it was possible to postulate a broad spectrum of different activation schemes, but not all. There does not appear to be an elegant method to implement synchronous activation, in which agents' future states are stored as all agents decide, followed by simultaneous state-change. Only asynchronous activation was implemented.

4.1 Random Activation

There are several suggestions in the original paper that the authors chose asynchronous random activation. The initial papers on activation were published in the same year (1992) [3, 5], so it is not unexpected that Gode and Sunder would not consider elaborating on the issue.

In our instantiation, random activation merely means that traders are chosen at random from the set of all traders. These traders form an offer. A turn is defined as complete when a number of traders equal to the total number of traders has made an offer. No data points are collected at the end of one turn, and no values are reset. All offers to sell or buy that are in the auction at the end of a turn continue in force at the beginning of the next turn. In fact, these offers are frequently canceled. The original model design had all offers canceled once a trade was complete.

Initialization and reinitialization: On the first activation, and every time the offers have been canceled, the first trader's offer will establish the new “best offer” of that type. Thus, if a seller is chosen first, he will choose a proposed sell price that is a uniform random variable between zero and his cost (for this item in his inventory sequence). A buyer will, likewise, establish the new “best buy” offer. Trading can commence as early as the second offer.

4.2 Uniform Activation

Asynchronous uniform activation is executed in a manner similar to random activation. At the beginning of each turn, the array of traders is shuffled. In one turn of uniform activation, all traders will be activated. Otherwise, the trade rules are the same: offers are carried over from turn to turn, but are canceled once a trade is complete. Initialization and reinitialization are conducted in the same manner.

The trade timing plots for market 3 are shown for the uniform activation scheme. There does not appear to be any significant difference in trade timing behavior between random and uniform.

4.3 Poisson Activation

Poisson activation is a process in which agents are activated according to an exponential distribution with an arrival rate, λ_A . This will mean that activations for any given agent are a Poisson process. In its simplest form, a Poisson activation scheme would have all agents activated with the same λ . This, however, would merely replicate the random selection method so we explore only the case of heterogeneous values for λ_A .

Poisson activation differs from other asynchronous methods in that this variation among the agents can be based on the state of

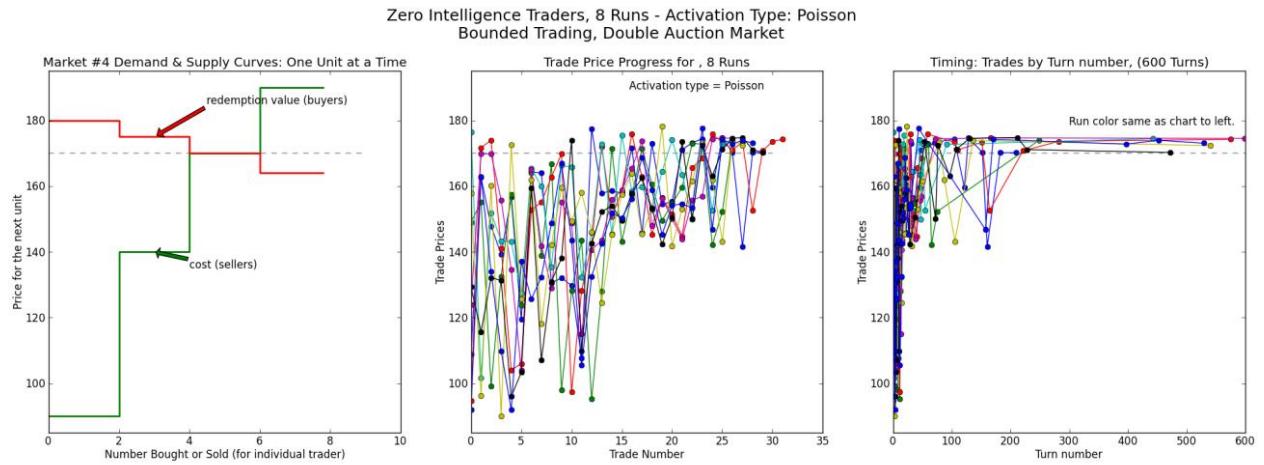


Figure 4. Market 4 (Poisson) Trade Price vs. Trade and vs. Turn

each agent or some internal parameter value. For our explorations, we chose agent wealth, which was calculated at the beginning of each turn. Thus, agent activation rates are made proportional to agent wealth values. In order to investigate the ‘leveling’ nature of these computer-based trading markets – a key question for the original researchers – we chose to make activation rates proportional to the absolute distance between the agent’s wealth and the average wealth of the population of agents. In that way, agents that are at the extremes (rich or poor) will likely trade more often.

In order to make appropriate comparisons between Poisson activation and other activation methods, it is necessary to re-normalize all of the values of λ_A so that, on average, each turn there will be one full population of traders’ activations. we accomplish this by building activation time for each agent and adding it to an ‘event list’. Trader-agent activation times are drawn sequentially from an exponential distribution and each added to the previous until the times exceed 1.0. These times are then all sorted and the trader agent sequence that results from that is passed to the program as a list of activations. Offer-making proceeds in accordance with this list for a given turn. At the beginning of the next turn the values of λ_A are again calculated and another sequence is generated. The order of each turn’s sequence is dependent on the current values of trader wealth and on a random draw.

This process works well once the model is established, but at the beginning of the model no trades have taken place and, thus, traders have no wealth. In these cases the values of λ_A are merely assigned randomly (and normalized as above). Once one trader has acquired some wealth the process can proceed as designed.

The Poisson process takes advantage of the ‘memoryless’ feature of the underlying exponential distribution. Thus, for every trader at the beginning of each turn can treat the ‘wait time’ as starting anew. It does not matter, given the waiting time is exponentially distributed, how long each trader has been waiting since the last activation.

4.4 Inverse Poisson Activation

The process of activating agents faster if they are further from the average has an interesting counterpart: activation rates that favor proximity to the average. Thus, we examined a λ -setting process that slows down agent activations when the trader wealth is farther

from the mean wealth. This inverse Poisson activation rate is the fourth activation scheme to be examined in the four markets.

It is important to note that the two Poisson schemes represent a conceptual departure from the other two asynchronous schemes. In varying the activation rate based on the agent state, we are examining *endogenous* activation. At least one article [1] has found that this can show differences in outcome behavior when compared with the more normal exogenous activation.

5. OUTCOME BEHAVIOR METRICS

Gode and Sunder do not rely heavily on precise quantification of the market results. This is consistent with their goal of measuring the performance of an automated market against that of a human market. They are trying to determine how much market efficiency (in profit creation and distribution) is due to the constraints of profit and loss rules and how much is due to human trading. Thus, they take the unconstrained automated market and the human market as two extremes and see where the bounded ZIT market falls. They judge that it falls much closer to the human market, but this is generally a qualitative judgment.

We chose to measure three aspects of the constrained ZIT market: its efficiency in generating wealth (or profits), its effectiveness in evenly allocating wealth among the traders, and the time it takes to reach equilibrium. Gode and Sunder used the first two measures in their paper, but left the third unexamined.

5.1 Wealth Generation

It is a straightforward matter to measure total wealth at the end of a run. One of the key (and unstated) influences on this total is the length of a run. Gode and Sunder ran a trading ‘day’ for 30 seconds. In our runs, we made use of the turn structure to better standardize the runs, choosing 5000 turns as a standard run.

The total wealth in the market is compared with the total *theoretical* wealth. Smith’s definition of market efficiency was used [6]. Thus, the allocative efficiency of a market is the total profits earned in one run (added across all traders at the end of the run) divided by the maximum profits available. Actual human markets quickly converge to 99% efficiency. Markets only vary from this, the authors noted in 1992, when typographic errors in market orders create a distortion in the price time series. (Considering the events of the past two decades, the Gode and Sunder paper could be seen as an important early warning of such market ‘errors’.)

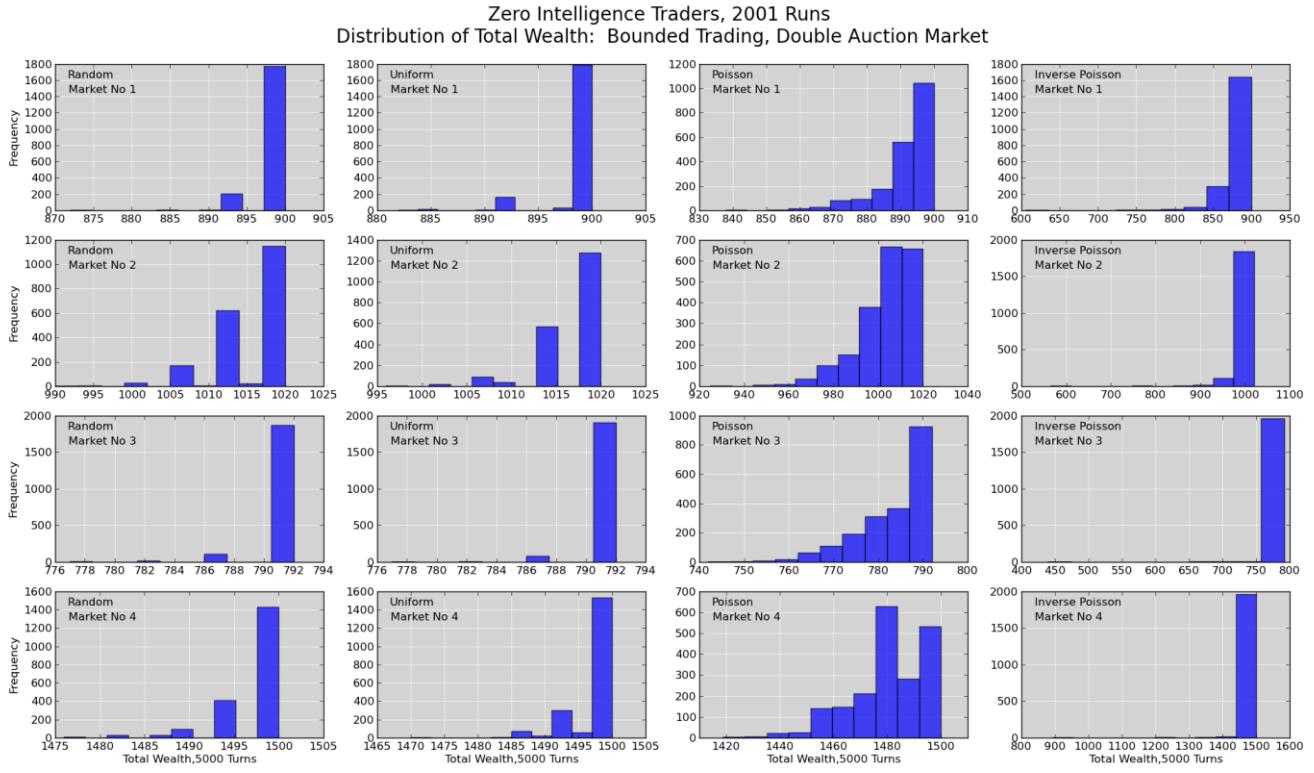


Figure 5. Total Wealth (All Traders) After 5000 Turns – Variable Scale

5.2 Profit Allocation

The second metric chosen by Gode and Sunder was the profit allocation among the traders. To determine this, they calculated the cross-sectional root mean squared difference between the actual and the equilibrium profits across the traders. They defined the value a_i as the profits (or total wealth) acquired by trader i . They also calculated the theoretical profits for this trader as π_i . Thus, the dispersion across all traders becomes

$$D = \sqrt{\frac{1}{n} \sum_i (a_i - \pi_i)^2} \quad (1)$$

They left unstated how they calculated the equilibrium values. We divided equilibrium profits into those for buyers and those for sellers. We assumed buyers' equilibrium profits as the profits they could earn if they traded all the shares they could at the market clearing price. This, of course, would only include those shares with a redemption value above the market clearing price. Similarly, the sellers' values of π_i was determined as the profits a seller would earn if all those shares held with costs below the market clearing price were sold at the market clearing price. Thus, to calculate D , it is necessary to separate the calculation of the sum into two parts. More correctly, it should be:

$$D = \sqrt{\frac{1}{n} \left[\sum_s (a_s - \pi_s)^2 + \sum_b (a_b - \pi_b)^2 \right]} \quad (2)$$

Where $s = \text{seller } s \in S$ and $b = \text{buyer } b \in B$ and $n = \text{the total number of traders}$. This separation is necessary because the supply and demand curves are not symmetrical. Sellers' equilibrium profits differ from those of buyers in essentially all markets.

5.3 Time to Last Trade

Gode and Sunder did not examine the model behavior over the long term for a variety of reasons. They were comparing simulated markets with actual human experiments. The human experiments had a finite duration because they were limited by many factors that are not present in simulations. Thus, the simulated markets were truncated and the long-term data are missing (or, in the terminology of statistics, the data were 'censored').

We expected to run the markets to exhaustion. That is, we experimented with a number of lengths of runs in the random and uniform activation types to find a reasonable point at which trading ended. We chose a run length of 5000 turns, believing this would encompass all trades for all markets and all activations. As noted in the result section, there was still censored data even at these extended runs. In fact, this represents a major difference among the activation schemes. Thus, while we didn't collect a comprehensive set of data, analysis of the turn at which the 'last trade' took place certainly achieved one of the key goals of this project – differentiating among activation schemes.

6. MODEL RESULTS

A full spectrum of experiments was run: four activation schemes across four markets. Each experiment consisted of 2000 runs of the market and activation, with each run including 5000 turns. At the end of each run, total wealth, wealth dispersion, and the turn of the last trade were recorded.

Market analysis shows that the exogenous activation schemes run to completion and the endogenous schemes (the Poisson activation types) still have some trading opportunities available at the end of 5000 turns. This is most apparent in the results in Figure 7.

Zero Intelligence Traders, 2001 Runs
Histogram of RMS Wealth Distributions, Market 3
Bounded Trading, Double Auction Market

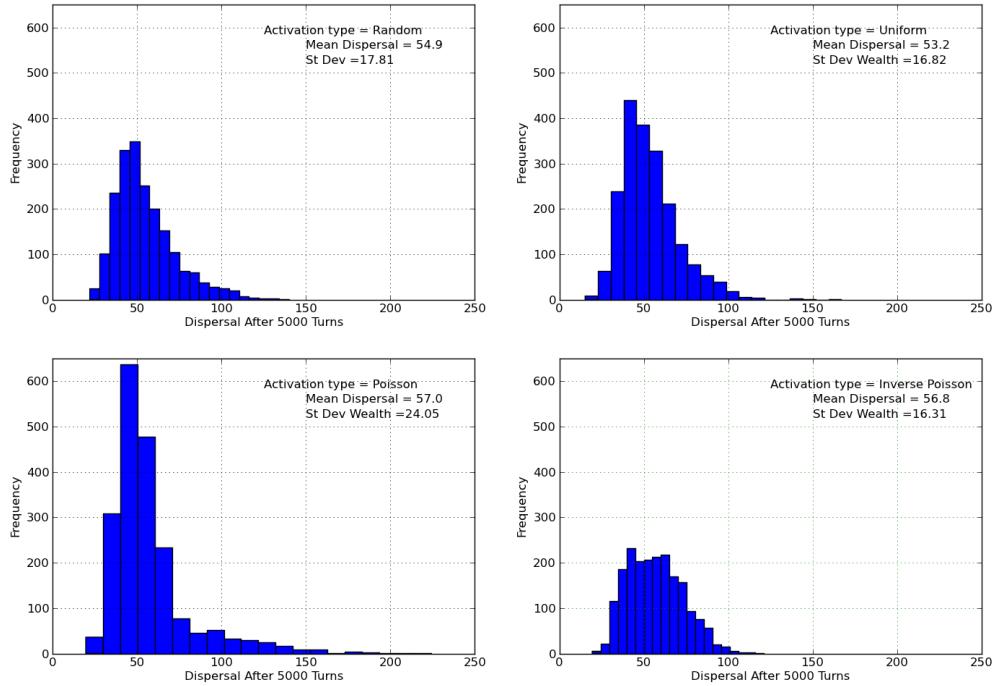


Figure 6. Wealth Dispersal, Market 3 (Constant Scale)

Figure 5 shows the four histograms of total wealth for all markets. The inverse Poisson activation exhibits extreme values of low wealth, but actually bunches much of the wealth closer to the maximum value for each market.

markets – allow the rejection of the null hypothesis for markets 2 and 4. Note that values that are too small to calculate are reported as 0. (While the averages are close, the power of the test is derived from the $n = 4000$ combined data points for the pair.)

Table 1. Mean Total Wealth at End of Run

| Activation | Average Total Wealth | | | |
|-----------------|----------------------|--------|-------|--------|
| | 1 | 2 | 3 | 4 |
| Random | 899.0 | 1016.6 | 791.6 | 1497.7 |
| Uniform | 899.2 | 1017.2 | 791.7 | 1498.2 |
| Poisson | 892.2 | 1003.2 | 785.4 | 1480.5 |
| Inverse Poisson | 881.3 | 999.6 | 785.1 | 1488.8 |
| Max Wealth | 900 | 1020 | 792 | 1500 |

With 2000 runs, it is possible to test the hypothesis that these means are drawn from different populations against the null hypothesis that the variation is simply due to random errors (and that the random errors are normally distributed).

With four activation schemes there would be sixteen pairwise comparisons. It is not necessary to examine these exhaustively to see differences among the activation types. As Table 2 shows, most of these comparisons are highly significant. Even the random-uniform comparisons – the closest averages for all the

Table 2. *p*-values for Average Total Wealth

| Comparison | Market | | | |
|-------------------|-------------------------------|----------------------------------------|-------------------------------|----------------------------------------|
| | 1 | 2 | 3 | 4 |
| Random - Uniform | 0.021 | 1.3×10^{-5} | 0.035 | 4.1×10^{-5} |
| Random - Poisson | 10^{-200} | 0 | 10^{-213} | 0 |
| Random - Inverse | 10^{-242} | 10^{-124} | 10^{-42} | 10^{-31} |
| Poisson - Inverse | 10^{-102} | 1.5×10^{-7} | 0.525 | 10^{-28} |

Gode and Sunder compared the total wealth in the simulated markets to the maximum total wealth possible. This maximum is shown on the final row of the wealth table for each of the four markets. Their objective was to compare how close the simulation came to maximum wealth with the proximity of the human markets. They deemed that their simulations across the four markets achieved essentially the same results as the human market, with efficiency percentages between 96 and 98%. These results were replicated in all markets by all activation types. The lowest percentage was 97.9% in the case of the inverse Poisson in Market 1.

Zero Intelligence Traders, 2001 Runs
Turn Number of Last Trade: Bounded Trading, Double Auction Market

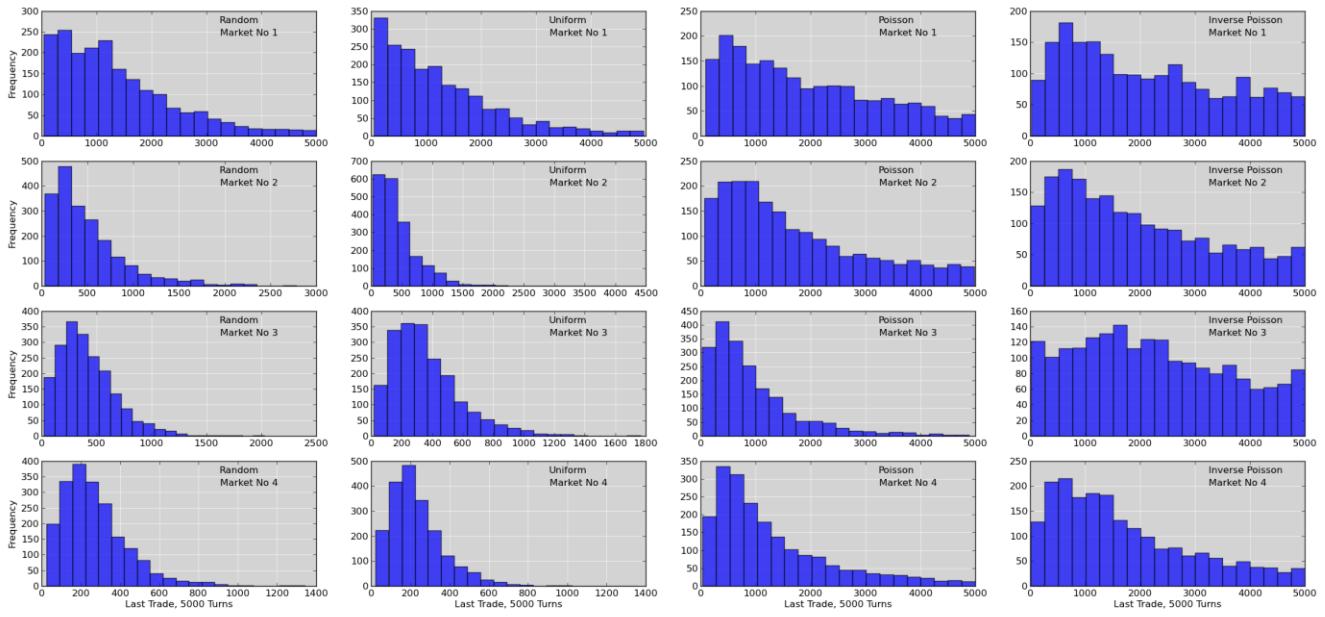


Figure 7. Last Trades in 5000 Turns – Variable Scale

Similar analysis can be conducted on the much more bell-shaped wealth dispersion. Wealth dispersion is depicted on the histograms on Figure 6. These have all been adjusted so that they appear on the same x - and y -axis scales, which we designate with a white background. With the scales adjusted, it's clear that the histograms appear significantly different. The Poisson activation histogram shows a significantly larger tail than the others. This may not be apparent from the small size of the bars on the far right hand side of that plot, but the automatic adjustment of the graphing program clearly adjusts for larger bins for the Poisson case to accommodate the larger range of data.

Table 3. Mean Wealth Dispersion, All Markets, All Models

| Activation | Average Wealth Dispersion | | | |
|-----------------|---------------------------|------|------|-------|
| | 1 | 2 | 3 | 4 |
| Random | 29.2 | 51.2 | 54.9 | 110.3 |
| Uniform | 28.8 | 50.5 | 53.2 | 111.0 |
| Poisson | 31.6 | 51.9 | 57.0 | 102.9 |
| Inverse Poisson | 28.7 | 51.8 | 56.8 | 110.9 |

While the wealth dispersion appeared to vary little across the runs, the large number of runs allowed us to determine that many of these differences were statistically significant. Using similar calculations to the averages of the wealth, we can develop another table of p -values. In this case, somewhat fewer of the pairings show differences that are significant. Market 3 shows some interesting behavior in that even the random – uniform comparison results in a difference that is significant at the 99% confidence level. Still, we reject the null hypothesis that the differences

between these sample means is a product of random fluctuations in seven of the 16 cases examined. Activation type makes a difference, at least statistically.

In addition to the odd shape of the Poisson activation histogram, it's also clear that the inverse Poisson activation type has a much tighter bunch of averages. The means between the two are quite similar (57 and 56.8), but the standard deviation is substantially larger for the Poisson activation scheme. In fact, the inverse Poisson standard deviation is nearly equal to that of the exogenous activation types (random and uniform).

Table 4. p -values Wealth Dispersion

| p -values | Market | | | | |
|---------------------------|------------|-------------------------|-------|---------------|-------------------------|
| | Comparison | 1 | 2 | 3 | 4 |
| Random - Uniform | | 0.23 | 0.20 | 0.002 | 0.36 |
| Random - Poisson | | 10⁻¹² | 0.278 | 0 | 10⁻²⁷ |
| Random – Inverse Poisson | | 0.14 | 0.28 | 0.0004 | 0.457 |
| Poisson - Inverse Poisson | | 10⁻²⁵ | 0.59 | 0.67 | 10⁻⁴⁹ |

Finally, we analyzed the evolution these markets and activation schemes over the long term. Gode and Sunder did not consider the dynamics of their simulation during extended runs because they were comparing them with human traders in finite-time markets. We recorded the turn at which the last trade took place before the end of run and use this as a metric for market closure. In evaluating the results, it appears that 5000 turns was more than adequate for the random and uniform activation methods, but that

Poisson and inverse Poisson were still exhibiting trading behavior late during a 5000-turn run (!).

Figure 7 shows the behavior of all four last trades for the four activation schemes. Clearly, for all markets, the extent of the trading varies substantially as the activation type is changed. Not only are the histograms of somewhat different shape, the Poisson and inverse Poisson clearly have censored trading activity.

Table 5. Mean Last-Turn Over 2000 Runs

| Mean Turn of Last Trade 5000-Turn Experiment | | Market | | |
|-------------------------------------------------|--------|--------|--------|--------|
| Activation | 1 | 2 | 3 | 4 |
| Random | 1377.2 | 503.5 | 415.1 | 270.0 |
| Uniform | 1273.4 | 438.3 | 357.6 | 234.2 |
| Poisson | 1919.3 | 1718.9 | 947.4 | 1300.1 |
| Inverse Poisson | 2124.9 | 1927.7 | 2240.6 | 1695.1 |

This phenomenon would affect analysis of any ZIT models, especially if trading were cut off after a few hundred turns. It is uncertain where Gode and Sunder stopped trading. They set their cutoff at 30 seconds of computer time, which itself might be a different measure for endogenous than for exogenous activation. In executing our simulations, the random and uniform experiments take about half the time as the two Poisson activation experiments.

Table 5 shows a full factorial analysis of the actual values of the mean. The sizeable difference can be observed by inspection, but a complete analysis of the p -values confirms the statistical significance of the result. There is no pairing that has a p -value larger than 5×10^{-11} . Thus, it can be concluded that activation makes a potent difference in the later stages of the ZIT model.

7. CONCLUSIONS

There are several motivations behind the question: Does activation change the outcome of agent-based models? Our simulation appears to answer different questions in different ways.

For the simple issue of analyzing statistical results, the analysis shows that for all three metrics (total wealth, wealth dispersion, and the last-trade parameter), there are statistically significant differences between at least some of the activation schemes, and for one metric there are significant differences among all of them.

We chose a ‘real world’ model – as opposed to a model of abstract agents engaged in mathematical game theory – to observe the impact of activation differences on policy recommendations. Gode and Sunder wanted to determine whether markets are made efficient by structural features (such as the requirement to make profitable trades) or by the rational decisions of human traders. They determined, using qualitative (but quite reasonable) analysis, that the constrained ZIT simulation essentially replicated the efficiency of the human traders in achieving the total theoretical wealth. They also concluded that simulated traders distributed the wealth close to but a little more than the human traders, at least in the early stages of trading. After a time, the human traders dispersed their profits more evenly, but this was undoubtedly due

to the memory effect. Simulated traders forgot their supply and demand curves at the beginning of each experiment.

Would Gode and Sunder’s conclusions have been different if they used different activation schemes? Probably not:

- All activation schemes and all markets ended with a total wealth that was between 97.92 and 99.96% of maximum wealth.
- Profit dispersion has a somewhat higher variance for the endogenous activation patterns, so it is possible that, given that they only did six runs, the authors might have generated outlier results. If they increased the number of runs, however, they would have returned to their original conclusion (simulated ZIT traders produce slightly larger dispersion, but far closer to human traders than unconstrained trading).

Gode and Sunder did not examine the question of model convergence or trade evolution. Thus, they would not have noticed the significant differences that appear in the last-trade statistics among the different activation schemes.

A third motivation for evaluating the importance of activations schemes is in establishing a proper standard for research in which the agent-based models of one scientific team are replicated by subsequent researchers. The Gode and Sunder article was chosen because it appeared as a reference in 1171 subsequent articles. Clearly, many other researchers are at least working with the concept of simulating markets, and many are actually building agent-based models using the zero-intelligence trading paradigm. (None of those 1171 use the words “Updating” or “Activation” – or their derivatives – in the title, so activation is not a major research focus in this domain.) In the research reported above, the differential results from last trade analysis alone (if not all the results) show that if a replication of ZIT model is expanded beyond the work of Gode and Sunder, the results must be shown to be robust over different activation schemes. Thus, if agent-based researchers are to meet the standard of other sciences and work on replicating one another’s experimental results, then reports of their results must include the activation scheme used in the model.

8. REFERENCES

- [1] Fernández-Gracia, J. et al. 2011. Update rules and interevent time distributions: Slow ordering versus no ordering in the voter model. *Physical Review E*. 84, 1 (Jul. 2011), 015103.
- [2] Gode, D.K. and 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 (Feb. 1993), 119–137.
- [3] Huberman, B. and Glance, N. 1993. Evolutionary Games and Computer-Simulations. *Proceedings of the National Academy of Sciences of the United States of America*. 90, 16 (Aug. 1993), 7716–7718.
- [4] Mitchell, M. 2009. *Complexity: A Guided Tour*. Oxford University Press, USA.
- [5] Nowak, M. and May, R. 1992. Evolutionary Games and Spatial Chaos. *Nature*. 359, 6398 (Oct. 1992), 826–829.
- [6] Smith, V.L. 1962. An Experimental Study of Competitive Market Behavior. *Journal of Political Economy*. 70, 2 (Apr. 1962), 111–137.