Extended abstract (~1,000 words), submitted to The First Conference on Zero/Minimal Intelligence Agents.
October 22-24, 2020; Yale School of Management, New Haven, CT, USA.
https://som.yale.edu/event/2020/10/the-first-conference-on-zerominimal-intelligence-agents

Methodological Mess-ups in Modelling Markets with Minimal-Intelligence Agents

Dave Cliff
Department of Computer Science, University of Bristol, Bristol BS8 1UB, UK.
csdtc@bristol.ac.uk, +44 79 77 55 22 50

This paper reviews past work on zero- and minimal-intelligence trader-agents, integrates it with extensive brand new results from a contemporary public-domain open-source market simulator, and argues (with the benefit of 20:20 hindsight) that we probably need to start again because of problematic methodological issues in previous work.

The review of past work presented here plots a progression starting with Gode & Sunder's seminal ZIU/ZIC agents[1], and Kaplan's Sniper[2] agent, through the well-known ZIP[3], GD[4], MGD[5], and GDX[6] trader-agent strategies, and ending at Vytelingum's AA[7] which until very recently was widely held to be the "best-performing" minimal-intelligence trader-agent strategy in the public domain. At each step in that progression, there was often an implication or explicit claim that the latest trader-agent was an improvement on those that had been published previously, and so a dominance hierarchy or "pecking order" was established among this set of trader-agent strategies. In this paper I show that the various researchers involved in this progression collectively converged on a somewhat restricted style of experiments, of methodology, for evaluating and comparing the various agent strategies, much of which can be traced back directly to Vernon Smith's ground-breaking experiments published in his original 1962 JPE paper[8], and the rest of which can be charitably attributed to the researchers not having convenient or affordable access to supercomputing facilities, which limited the number of computer simulation experiments that they could perform. The new results in this paper arise from releasing these restrictions: I summarise here results from several million simulated market sessions (where one simulated *market session* is equivalent to a single market experiment in Smith's seminal work) and in which the market sessions relax many of the simplifying assumptions introduced by Smith. The results in this paper come from millions of market simulation sessions in which key parameters that were largely held fixed in previous work are systematically varied across their ranges; and in which for each individual market session several of the Smith-style simplifying assumptions that had previously been commonplace to research in this field are relaxed or abandoned. And, when these changes are made and the consequences are explored exhaustively, we find that the old pecking order no longer always holds true: in some cases, simple zero-intelligence trader-agent strategies such as ZIC can outperform more complex adaptive strategies.

Methodological hangovers from Smith's early experiments include: (1) the use of regular periodic synchronous re-distribution of allocations of stock and cash to all traders in the market at the same time; (2) the use of constant or only occasionally changing supply and demand schedules that lead to the model market's competitive equilibrium price (denoted here by P0) remaining constant for protracted periods within each market experiment; (3) simulated market mechanisms based on "open outcry" trading pits (where any trader can announce a quote at any time, and only the most recent bid or offer is acted upon) rather than a Limit Order Book (LOB) and (4) a focus on summary statistics such as Smith's alpha, which measure convergence of transaction prices toward the theoretical P0 value.

Moreover, in publications where comparisons were made between two trader-agent strategies, this often involved A/B testing via one or both of two methods: either *one-in-many* (OIM) tests; or *balanced group* (BG) tests. In OIM tests, a single trader-agent of type B was evaluated when operating within a market that was otherwise homogeneously populated by agents of type A; whereas in BG tests, the market was populated with a 50:50 ratio of agent-types A and B. If the market is set up to have an even number N of trader-agents, then we can characterise BG tests as involving a A:B ratio of $(^N/_2):(^N/_2)$, at the midpoint of the range of possible ratios; and OIM tests as involve an A:B ratio of (N-1):1, at an extreme endpoint on the range of possible ratios. Curiously, often no other ratios were used in evaluation and

---

[1] Gode, D. & 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.

[2] Rust, J., Miller, J., Palmer, R., 1992. Behavior of Trading Automata in a CDA Market. In Friedman, D., Rust, J. (eds) *The Double Auction Market: Theories and Evidence.* Addison-Wesley, pp.155-198.

[3] Cliff, D., 1997. *Minimal-Intelligence Agents for Bargaining Behaviours in Market-Based Environments.* Tech. Report HPL-97-91, Hewlett-Packard Labs.

[4] Gjerstad, S., Dickhaut, J., 1997. Price Formation in Continuous Double Auctions. *Games & Economic Behavior*. 22(1):1-29.

[5] Tesauro, G., Das, R., 2001. High-Performance Bidding Agents for the Continuous Double Auction. *Proc. 3rd ACM Conf. on E-Commerce*, pp.206-209

[6] Tesauro, G., Bredin, J., 2002. Sequential Strategic Bidding in Auctions using Dynamic Programming. *Proceedings AAMAS2002*.

[7] Vytelingum, P. *et al.* 2008. Strategic Bidding in Continuous Double Auctions. *Artificial Intelligence,* 172(14):1700-1729.

[8] Smith, V., 1962. An Experimental Study of Competitive Market Behavior. *Journal of Political Economy* 70(2):111-137.

comparison of the various trader-agent strategies, and one very recent result[9] that we review and expand upon here is that whether A outperforms B or vice versa is often dependent upon the A:B ratio, and testing only at the endpoint or the midpoint of the range of possible ratios can give misleading results that lead to incorrect conclusions about the dominance hierarchy.

All the new experiments reported in this paper have been conducted on the *Bristol Stock Exchange* platform,[10] a freely available open-source simulation of a Limit-Order-Book (LOB)-based financial exchange, for a single tradeable asset, written in Python, which has been in widespread use since its first release in 2012. Large-scale experiments were conducted using a parallel processing architecture running in the cloud via Amazon Web Services. At the time this paper is presented at the Yale conference in October, the source-code and sample data will be made freely available.

Furthermore, results shown for the first time in this paper demonstrate that introduction of new deterministic zero-intelligence strategies can have a disruptive effect on the dominance hierarchy. These two new strategies are named SHVR and GVWY: SHVR is a "shaver" strategy that simply aims to always beat the current best bid or offer, "shaving" the current best price by one penny; and GVWY is a "giveaway" strategy that simply issues quote prices (bids or offers) at the limit price specified in that trader's individual allocation, i.e. GVWY seeks to extract zero surplus on each transaction. Depending on what trader-agent strategies are present in the market, and at what ratios, both SHVR and GVWY can be surprisingly profitable (i.e., can extract surprisingly large amounts of the available surplus). The result that GVWY can be consistently profitable, despite its manifestly non-profit-making hardwired strategy, is somewhat counterintuitive but we demonstrate that it is readily explainable as a consequence of the traders' price-discovery interactions mediated by the "matching engine" within the LOB-based market mechanism used in BSE, which is the same matching engine found in all major real-world electronic financial exchanges.

In summary, this paper summarises results from millions of simulated market sessions in which the primary findings/contributions are:

- Demonstration that the dominance hierarchy among well-known trading-agent strategies (ZIC, ZIP, GD/MGD/GDX, and AA) is affected by the extent to which various simplifying assumptions, common in the previous literature, are true (or not) in the market that the agents are tested in. In essence, if you test these agents in the simple minimal market-model environments that they were first developed for, you get the dominance hierarchy that is recorded in the literature; but once you relax those simplifications, making the market environment closer real-world markets, the dominance relationships may alter in surprising ways.
- Demonstration that two deterministic zero-intelligence trader-agent strategies, SHVR and GVWY, can be surprisingly profitable, as a consequence of the order-matching system commonplace in any real-world limit-order-book automated financial exchange.

[9] See Cliff, D., 2019. Exhaustive Testing of Trader-Agents in Realistically Dynamic CDA Markets. In *Proc. ICAART-2019*; and Snashall, D., Cliff, D., 2019. Adaptive-Aggressive Traders Don't Dominate. In van den Herik, J., Rocha, A., Steels, L., (eds) *Agents & Artificial Intelligence,* Springer.

[10] *Bristol Stock Exchange*. GitHub source-code repository: https://github.com/davecliff/BristolStockExchange, 2012. See also: D. Cliff. BSE: A Minimal Simulation of a Limit-Order-Book Stock Exchange. *Proceedings of the European Modelling and Simulation Symposium*, pages 194-203, 2018.