DEPARTMENT OF ECONOMICS

ISSN 1441-5429

DISCUSSION PAPER 12/17

The Double Auction Market and Inequality

Paul Brewer[1] and Anmol Ratan[2]

Abstract:

Do markets exacerbate inequalities? We compare distributional outcomes produced by robot-populated double auction markets to a well-defined theoretical benchmark. We examine robot trading in two double auction environments with identical aggregate supply and demand curves but different individual agent supply and demand curves. The Law of Supply and Demand predicts the same theoretical competitive equilibrium price and quantity for the two treatments. The first market environment is constructed such that in competitive equilibrium the resulting profits are equal for all agents. The second treatment is constructed to yield substantial income inequality in competitive equilibrium. We model two types of trading strategies- "zero intelligence" traders and "snipers" who are more sophisticated than "zero-intelligence" traders in two market environments. We compare our observations to the outcomes predicted by competitive equilibrium theory regarding transaction prices, allocative efficiency and the distribution of profits among traders, varying the proportion of snipers in the market. The predictions of the competitive theory are consistent with outcomes for low volumes of snipers in the market. As the volume of snipers in the markets increases, the competitive predictions, however, are not consistent with observed outcomes. Most importantly, in scenarios where allocative efficiency falls short of competitive equilibrium predictions, the inequality based on the observed distribution of profits among traders is larger than the inequality predicted by the competitive equilibrium.

Keywords: Markets: Double auctions; competitive equilibrium; efficiency; inequality; experiments

Acknowledgements: The financial support for this research was provided by Monash University Department of Economics Research grant. Access to numerical simulation software running on Google's on-demand compute engine was provided by Economic and Financial Technology Consulting LLC.

[1] Paul Brewer; Economic and Financial Technology Consulting LLC Email: drpaulbrewer@eaftc.com

[2] Anmol Ratan; Monash University Email: anmol.ratan@monash.edu

The Double Auction Market and Inequality

1. Introduction

"*Although average economic well-being has increased considerably over time, the degree of inequality in economic outcomes has increased as well*" (Bernanke 2007). Due to increasing social concern and awareness of rising inequalities, a tradeoff facing economists in choosing between market-based incentives to spur growth and efficiency is that these same incentives may also contribute towards greater inequality. Bernanke (2007) summed up this situation and stated that: "*No objective means of answering these questions (regarding distributional outcomes produced by markets) exist. One can only try to understand the various issues and tradeoffs involved and then come to a normative judgment based on that understanding*." Many political and moral leaders tend to hold markets responsible for rising inequalities, as captured in the following statement by Pope Francis (2015): *"…as long as the problems of the poor are not radically resolved by rejecting the absolute autonomy of markets and financial speculation and by attacking the structural causes of inequality, no solution will be found for the world's problems…."*[3] Others praise market capitalism for lifting the world from poverty and creating greater equality, with claims such as "*700 million people lifted from poverty since 1990,*" greater access by the poor to products and services once reserved to the rich and lengthening and reduction in inequality of the human life-span (Delsol 2017). Piketty (1997, 2014) caricatures this debate as almost purely a self-interested, political conflict that can be addressed by accumulation and examination of historical facts and trends. In this paper, we take a different approach. Although a variety of causes including skill-based technical changes, increased globalization, institutional changes, and tax policy, have been suggested as responsible for rising inequality in recent times,[4] there may be a lack of understanding regarding the role of markets themselves towards exacerbating the inequalities. We believe this role of markets can be isolated from the complex milieu of world events and studied using experimental methodologies.

The following general questions motivate our investigation of the distributional properties of markets:

1. Do market processes worsen inequality or merely duplicate inequality already inherent in the economic environment?
2. Given that the neoclassical theory does so well in predicting prices and allocative efficiency in market environments with imperfect information and little or no regulation of prices, does it succeed to explain the distribution of incomes?

[3] Cited in Collins (2017)

[4] The potential causes of unequal outcomes and rising inequalities has been studied with great interest in the development literature. The primary causes identified in the case of underdeveloped economies are slightly different from those for developed economies. *"…while specific grievances varied from country to country and, in particular, that the political grievances in the Middle east were very different from those in the West, there were some shared themes. There was a common understanding that in many ways the economic and political system had failed and that both were fundamentally unfair* (pg. ix, Stiglitz 2012)."

3. Can we learn something about the inequality in the real world by observing the outcomes generated by market processes that do not discriminate between agents except for their cost and value endowments? Is the inequality observed in the real world primarily produced by institutional features or market processes?
4. Is there some natural pattern that emerges in a macro sense but not in a micro sense (i.e., we all experience good and bad times, but there is no inherent bias)? Do markets systematically amplify the fortunes of the rich and hurt the poor? By contrast, do markets protect people?[5]

In recent times, economists have studied the historical trends in the world distributions of income and wealth (Stiglitz 2012; Piketty 1997, 2014) to identify various causes for rising inequalities. These approaches, however, have limitations that make it difficult to judge the relative contribution of various causes towards rising inequalities. First, although historical comparisons are possible, the lack of a suitable benchmark against which distributional outcomes can be compared does not allow an unambiguous answer as to whether markets contribute to rising inequalities. Second, historical investigations do not typically allow detailed cross-sectional evaluations of welfare implications of market processes. Instead, historical investigations typically report large-scale aggregate statistics. To paraphrase Hayek (1945)[6], detailed information on everything that is going on in an economy is never in possession of a single person. The consequence is that the historical record, whether from the then-existing government or private efforts, is only a small fraction of the data contained in a functioning economy. Thus, further approximations are unavoidable for drawing inferences for the entire population in an economy.

Historical methodologies are not alone in suffering problems with incomplete data or approximation. Purely theoretical models can also over-simplify. Simple models of exchange proposed by physicists, not economists, predict the distribution of money in a suitably simplified closed system will evolve to substantial inequality in equilibrium. Dragulescu and Yakovenko (2000) repurpose the Boltzmann-Gibbs equilibrium distribution from statistical physics to economics. The result originally concerns the equilibrium of a closed system with many particles -- too many to track and analyze individually -- such as atoms of a gas trapped in a bottle and isolated from the outside world. Within the closed system, random zero-sum energy transfers occur between particles, but the total energy is constant. The result is repurposed to apply to a "closed economy" with random zero-sum monetary transactions between agents and constant total money. They predict that individual monetary holdings converge towards an

[5] As suggested by the "…Market as a Substitute for Individual Rationality…" title and conclusions of the well-known Gode and Sunder (1993) study of markets populated by "zero intelligence" robots.

[6] "… the data from which the economic calculus starts are never for the whole society "given" to a single mind which could work out the implications, and can never be so given. The peculiar character of the problem of a rational economic order is determined precisely by the fact that the knowledge of the circumstances of which we must make use never exists in concentrated or integrated form but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess." Hayek (1945), p.519.

exponential Boltzmann-Gibbs distribution. The model is easily replicated and demonstrated numerically by computer programs in work by these authors and others[7], and in at least one introductory agent-based modelling textbook (Wilensky and Rand, 2015). In these demonstrations, a computer evolves an "economy" of robot agents. Each robot has equal starting money that is updated over a sufficiently large number of interactions with randomly chosen agents, gaining or losing money in each interaction in a zero-sum manner such that the total money in the system is always constant. Over time, the system equilibrates towards the predicted exponential limiting distribution with its expected inequality. While interesting for generating an inequality result from very minimal assumptions, this class of models ignores well-established features of economies, including the following two that break the model's assumptions: (i) "trades" occur because both agents are better off, so positive gains occur from exchange instead of being zero-sum[8]; and (ii) the total amount of money in an economy is not constant. [9]

What are other alternatives to these approaches? We believe that methodologies from experimental economics can help plug some of the deficiencies of both the overly complex and overly simplified approaches to studying inequality. First, with suitable controls, we can derive theoretical benchmarks against which distributional outcomes can be compared; second, the welfare implications of markets that can be studied using these methodologies are observable in great detail; third, in simple controlled environments, we can focus on studying the distributional outcomes, which are produced by markets rather than not being able to interpret the outcomes which are often produced by the interaction of market and non-market factors; and, finally, it is straightforward to generate more data in response to questions, new theories, or other concerns.

To make headway into the well-recognized challenges posed by the broad questions, we begin by exploring the following simpler questions:

1. If we create an environment where the competitive theory, i.e., reasoning based on the Law of Supply and Demand, predicts equality of incomes at competitive equilibrium, do we observe outcomes consistent with this prediction?
2. If we create an environment where the competitive theory predicts unequal incomes, do we observe outcomes consistent with predicted inequality of incomes?
3. If we vary the sophistication of the strategies employed by traders, going from less sophistication to more sophistication, do we observe the same results?

[7] http://physics.umd.edu/~yakovenk/econophysics/ downloaded Feb 24, 2018.

[8] Besides upsetting the conservation (constancy) of total money, positive-sum trades also break an important time-symmetry property needed to simplify the Boltzmann equation. An elastic collision of ideal particles is still a valid elastic collision if time is played backwards, likewise a zero-sum trade is still a zero-sum trade if played backwards, but a positive-sum trade becomes a negative-sum trade if time is played backwards. According to Dragulescu and Yakovenko (2000), a violation of time-reversal symmetry would yield some other less simple stationary distribution or none at all (p. 727).

[9] It is well known that aspects of the money supply are continuously changing through common fractional-reserve banking practices and interventions by governments and central banks.

The answers to these questions can be useful at this stage. If an economic environment is purposely designed to yield equal outcomes in competitive equilibrium (henceforth CE), yet the market generates substantial inequality, then the reasons could be investigated. At the very least, the methodology we adopt in our paper can be utilized to compare the distributional properties of alternative market processes and may guide us to modifications of market rules that are more socially desirable. If market forces are shown to behave according to known principles, without amplifying the inequality inherent in an environment, then it can be inferred that market forces are not the cause because the cause was identified in the environment.

In the experimental literature, markets have been studied via the basic contracting institution of a double auction (henceforth DA), which preserves the law of demand and supply. We shall focus on this institution to analyze the distributional properties of markets given that it can capture unstructured decentralized interactions for the exchange of commodities. This approach has various advantages: First, the main predictions of CE regarding-price and allocative efficiency- are supported in various studies of DA markets with simulated traders and human participants, including some corner cases where the CE predicts zero profits for one side of the market (e.g., buyers) and positive profits for the other. Despite these results, the distributional outcomes of the DA have not been fully investigated with regards to the resulting level of inequalities.

Second, it is also, by now, well known that humans do not always maximize their own profit to the detriment of others in simple two-person Dictator and Ultimatum games. [10] When we place these general findings about individuals who refuse to be completely unfair when given the opportunity alongside the general findings about markets, it increases our uncertainty about the primary questions that drive our research. For example, it is not clear why individuals, who seem to display pro-social behavior in specific game interactions, display the opposite attitudes when they are competing with others on the same side of the market. Does this competition crowd out the pro-social tendencies in market environments? Given these, can we say anything about market allocation processes beyond their efficiency or appearance of converging prices that admits or refutes the critiques of self-interest by those studying inequality?

What do we accomplish in this paper? In this paper, we work with well-established methodologies to explore two market environments in which the CE predictions are identical for prices and allocative efficiency but vary in terms of ex-ante distribution of profits. We simulate trader behavior using two well-known trading strategies: zero-intelligence traders and snipers. It can be argued that these trading strategies vary in the level of sophistication since the zero-intelligence traders are merely budget constrained whereas the snipers seek higher profits in the marketplace. We analyze the outcomes that are likely to result from gradual infusion of snipers in the DA markets. This allows us to compare the outcomes with the benchmark environment with 100% zero-intelligence traders. Across two market environments the following results are

[10] For an overview of dictator and ultimatum games, we refer the interested reader to the relevant chapters contributed to Plott and Smith's *Handbook of Experimental Economic Results,* Chapters 46-50.

obtained: (i) with low volume of snipers in the markets, the average prices, efficiency levels, and distributional outcomes are consistent with the CE predictions; but (ii) in scenarios where the volume of snipers exceeds a threshold, the efficiency levels are much below the CE predictions, and the level of inequality based on distributional outcomes is larger than the CE predictions.

The predictions based on the CE regarding the distribution of profits are consistent with our results in a trading environment with 100% zero-intelligence traders. This is heartening to know that, eventually, no systematic biases in distributional consequences are observed. The market does not seem to amplify the inequality that is predicted due to the allocation of values and costs to the traders. When observing single trading periods, noisy trading, a term we will clarify later, does affect results but this noise seems to average out over a large number of periods. However, the observed outcomes with high volumes of snipers in the markets are not consistent with the CE predictions. This inconsistency suggests that a conclusion derived by Gode and Sunder (1993) -- regarding the secondary importance of human motivations and cognitive abilities with respect the institutional features of the marketplace -- does not generalize to scenarios where most traders are not zero-intelligence traders.[11]

The remainder of this paper is organized as follows: Section 2 provides some background and methodology used in relevant literature, Section 3 introduces the market mechanism (the DA market), market environments, predictions of the competitive theory (the Law of Supply and Demand) and trading strategies. Section 4 reports the results of our numerical experiments and section 5 and 6 conclude.

2. Background and Methodology

A DA market is a multilateral process in which buyers and sellers can freely submit bids or asks and accept asks or bids submitted by others. As discussed above, a DA market can capture unstructured decentralized interactions for the exchange of commodities. Participation in a DA market as a trader-buyer or seller-requires minimal training or experience since one-to-one bargaining frequently occurs in everyday life. Due to these advantages, major stock, currency, commodity, and other markets are organized as double auctions.

Earlier studies of behavior in DA markets assumed costless bargaining between traders, in a classroom or laboratory, who were either assigned the role of a buyer or a seller (Chamberlain 1948; Smith 1962, 1965; and Smith and Williams 1990). Wilson (1984) proposed the earliest known characterization of equilibrium in dynamic double auctions. Tests of individual behavior in DA markets produced some intriguing results. Early DA market experiments with human participants reported a surprising convergence towards the predictions of the competitive theory and the Law of Supply and Demand – even in environments with few traders and imperfect information. The aggregate predictions of CE regarding prices and efficiency were consistent with outcomes observed in a wide variety of settings. Even in extreme environments where

[11] They concluded that it is primarily the market institution that ensures this outcome, while "the effect of human motivations and cognitive abilities has a second-order magnitude at best" (Gode and Sunder, 1993, p. 133).

the competitive theory predicts that only one side of the market would earn all of the profit, both Smith (1965) and Smith and Williams (1990) found eventual, slow, price convergence to CE price in human-subject laboratory experiments. Experimental economists have interpreted these various observations to claim that markets work better than expected.[12] On the other hand, Cason and Friedman's (1993) reported that beyond the prediction of high ex-post trading efficiency in DA markets, individual behavior in laboratory experiments was inconsistent with the predictions of Wilson's model. To reconcile these results, other studies (Easley and Ledyard 1993, Friedman 1984, 1993; Giminez-Funes et al. 1998) suggested that rationality is not a necessary condition for observing efficient outcomes and convergence to competitive equilibria in DA market. The most striking confirmation of these observations was reported by Gode and Sunder (1993) who reported a series of numerical experiments that establish high efficiencies of allocation and a repeated price dynamic towards CE in DA markets populated by "zero intelligence" robot traders that simply bid or ask randomly within their respective budget constraints. They concluded by suggesting that "it is primarily the market institution that ensures this outcome, while "the effect of human motivations and cognitive abilities has a second-order magnitude at best."[13] Robot agents confined to such simple rule-of-thumb strategies may seem to be too simple but, together with market rules, can generate complex aggregate behavior and have produced practical recommendations for financial market policy.[14]

Wide interest in these results guided the quest to identify other explicit strategies and theories based on simple yet plausible rules of thumb that could be easily deployed in DA markets. Rust, Miller and Palmer (1993) reported a simple sniping strategy by Kaplan that outperformed other such strategies in a high stakes tournament. We will study the distributional consequences of this strategy in competition with the ZI Agents.

The construction of our experiments follows a causal framework ("Reiter's Triangle") described by Reiter (1977), later refined into experimental methodology by Plott(1982) and Smith (1982). It begins with the deconstruction of complex social processes into three separable units, modelled independently, that interact together to produce social outcomes: (1) the environment (also known as preferences, values, costs, endowments); (2) the institution (also known as processes or mechanisms, or more specifically, markets or auctions); and (3) the traders (including strategies, profit-maximizing or

[12] Summarizing Vernon Smith's contributions, Eckel, Houser and Boettke (2017) write: "He tells of his surprise when the first market [...altered from Chamberlin's setup to be more like a stock market...] converged to competitive equilibrium, and of his efforts to stress-test the environment to see if he could get a more reasonable result. However, as he might say, the darned thing continued to converge to equilibrium. Little did he know at the time that the double auction would prove to be the most powerful of market institutions, ensuring convergence with or without incentives, and with as few as three buyers and sellers." (pp. 639-640)

[13] These results, however, were not observed in alternative market environments such as those reported in Brewer, Huang, Nelson, and Plott (2002). In their experiments with continuously refreshed supply and demand, once the Marshallian path as a potential convergence dynamic is removed, price convergence is not observed.

[14] Brewer, Cvitanic and Plott (2013) use agents even simpler than the "Zero Intelligence" agents. Agents that simply mimic properties of an order flow together with Poisson arrival and various market rules allowed insight into the phenomenon of flash crashes, the relative effectiveness of corrective interventions, and the desirability of periodic "call markets" vs. tick-by-tick trading.

otherwise). To explore effects of the strategy component (3, above), we use computerized traders instead of human subjects. We shall work with *zero-intelligence* (hereafter, ZI) and *sniping* strategies to model trader behavior in DA markets in our paper. The ZI trading strategy environment, in which traders do not pursue profits and do not observe, remember, or learn but are restricted by market discipline, provides a useful benchmark with which outcomes could be compared when more sophisticated trading strategies may be employed.

Our analysis addresses the specific questions discussed above in a reduced environment, which allows us to implement controls that are necessary to achieve unambiguous predictions about the distribution of profits, based on the CE. Using simulated traders, besides conveying the usual advantages, allow studying the distributional consequences at the level of an individual unit. Such a level of detail contrasts with how distributional consequences have been studied thru summary metrics such as the Lorenz curves and Gini coefficients in the development literature.[15] Irrespective of these considerations, we shall explore the relevant metrics for the larger audience that may find our analysis pertinent.

3. Market Mechanism, Environments, and Competitive Predictions

Market Mechanism and Trading Rules

We utilize a DA market as our preferred environment to study the outcomes produced by market processes for the reasons discussed above. A DA market is a multilateral process in which buyers and sellers can freely submit bids or asks and accept asks or bids submitted by others.

3.1 Bid submission and trading rules

At any time in a trading period, a buyer could enter a bid by stating his or her desired price and quantity, or similarly, any seller could enter an ask by stating quantity and price. Trade occurs when the highest bid price exceeds the lowest ask price, and the price is typically set to the earlier of those two orders.

Several variations of the DA market are possible. The following features of our DA market design are similar to those in Gode and Sunder (1993): (i) Each bid, ask and transaction was valid for a single unit only, (ii) bids must be increasing, asks must be decreasing, with respect to the current best bid and best ask (iii) a trade occurs when a new bid is recorded which exceeds a pre-existing ask, or when a new ask is recorded that is less than a pre-existing bid. Unlike Gode and Sunder's DA market, in our design, unsuccessful bids and asks are not cancelled at each successful trade but retained from one trade to the next, until the end of the trading period. We do cancel bids and asks that are replaced by that same trader's improved bid or ask. When a trade occurs, the trading price is equal to that of the pre-existing bid or ask, whose acceptance is triggered automatically by the new entry.

[15] For example, distributional consequences studied via these metrices are identity neutral and need to be supplemented by other measures if welfare consequences were to be studied more carefully.

Besides these design features, each trader had an opportunity to submit their bid or ask that occurred randomly, at a constant probability per unit of "experiment time," resulting in arrivals of bids and asks following a Poisson distribution[16] with individual rate 0.1/sec (or approximately 1 bid/ask per 10 seconds per trader). Although the Poisson arrival distribution is used in other human-subject and numerical studies[17] the primary effects here are to (i) effectively shuffle trader arrival and eliminate any spurious phenomenon based on serial correlation that could occur when specific traders always meet in a pre-defined order; and (ii) provide a method to limit the number of opportunities to trade in a market period that is closer to the levels of trading activity that might be observed in human-populated markets while we are using robot traders that could otherwise attempt to trade millions of times per second.

3.2 Traders and induced values and costs

Each DA market consisted of 40 traders equally divided into 20 buyers and 20 sellers. At the beginning of each period, each buyer (seller) was assigned values (costs) for three units of an unspecified commodity "X" which could be redeemed upon successfully trading the unit(s). Each buyer (seller) is assumed to be privately aware of their own redemption value (cost) v_i (c_i) of each unit i, and the buyer's (seller's) profit from buying (selling) this unit at price p_i is given by $v_i - p_i$ $(p_i - c_i)$. Sellers had no fixed costs and incurred costs only for units sold. The distribution of redemption values v_i $(i = 1, 2, 3,$….. n) for units of "X" assigned to buyers, defined the aggregate demand schedule for the commodity. And the distribution of redemption costs c_i $(i = 1, 2, 3,$….. n) assigned to sellers, defined the aggregate supply schedule for the commodity.

Each market was run for a specified duration which allowed sufficient time for trading.

Market Environment: Demand and Supply Schedules

We explore two treatment conditions, named "Market 1" and "Market 2", such that:

(i) the aggregate demand and supply schedules are identical in each market; but
(ii) the distribution of buyers' values (sellers' costs) varies across buyers (sellers) systematically. The distribution of redemption values and sellers' costs for each market is described below.

[16] Technically, this is achieved by assigning time delays for future action from the conjugate exponential distribution.
[17] including Brewer, Cvitanic and Plott (2013); Alton and Plott (2007) and Farmer, Patelli and Zovko (2005).

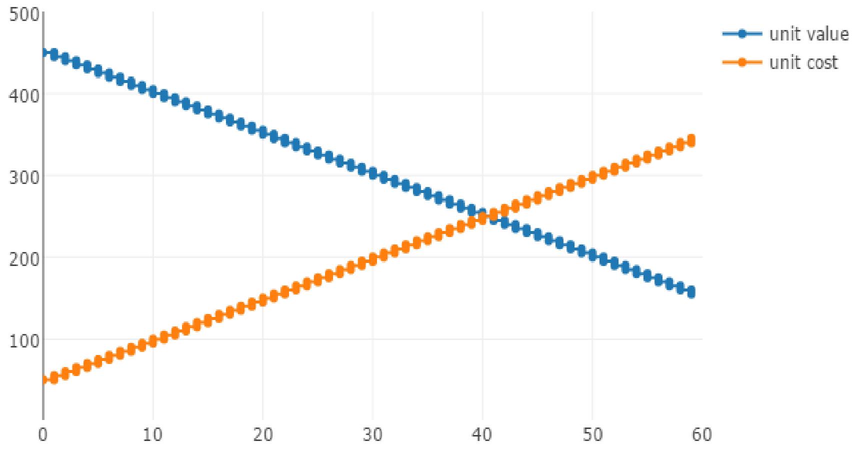

Figure 1: Aggregate Demand and Supply -- All Market Environments

Market 1

The schedule of induced values and costs in market 1 is shown in table 1:

Table 1: Values, Costs, CE Predicted Profits in Market 1
CE: Price=250, Quantity=40-41, Gini coefficient=0

Buyers					Sellers				
Id	*Values*			*Predicted profit*	*No.*	*Costs*			*Predicted profit*
1	450	255	250	205	1	50	245	250	205
2	445	260	245	205	2	55	240	255	205
3	440	265	240	205	3	60	235	260	205
4	435	270	235	205	4	65	230	265	205
5	430	275	230	205	5	70	225	270	205
6	425	280	225	205	6	75	220	275	205
7	420	285	220	205	7	80	215	280	205
8	415	290	215	205	8	85	210	285	205
9	410	295	210	205	9	90	205	290	205
10	405	300	205	205	10	95	200	295	205
11	400	305	200	205	11	100	195	300	205
12	395	310	195	205	12	105	190	305	205
13	390	315	190	205	13	110	185	310	205
14	385	320	185	205	14	115	180	315	205
15	380	325	180	205	15	120	175	320	205
16	375	330	175	205	16	125	170	325	205
17	370	335	170	205	17	130	165	330	205
18	365	340	165	205	18	135	160	335	205
19	360	345	160	205	19	140	155	340	205
20	355	350	155	205	20	145	150	345	205

Notes: (i) The CE price is 250 and (ii) the CE volume is 40-41 units. (iii) The frequency of inframarginal (extramarginal) buyers' values and sellers' costs are 40-41 units (19-20 units) respectively. Thus, the frequency of inframarginal and extramarginal units are the same across buyers' and sellers.'

Market 2

The schedule of induced values and costs in market 2 is shown in table 2:

Table 2: Values, Costs, CE Profit Predictions in Market 2

CE: Price=250, Quantity=40-41, Gini Coefficient = 0.166

Buyers					Sellers				
Id.	*Values*			*Predicted profit*	*No.*	*Costs*			*Predicted profit*
1	450	350	250	300	1	50	150	250	300
2	445	345	245	290	2	55	155	255	290
3	440	340	240	280	3	60	160	260	280
4	435	335	235	270	4	65	165	265	270
5	430	330	230	260	5	70	170	270	260
6	425	325	225	250	6	75	175	275	250
7	420	320	220	240	7	80	180	280	240
8	415	315	215	230	8	85	185	285	230
9	410	310	210	220	9	90	190	290	220
10	405	305	205	210	10	95	195	295	210
11	400	300	200	200	11	100	200	300	200
12	395	295	195	190	12	105	205	305	190
13	390	290	190	180	13	110	210	310	180
14	385	285	185	170	14	115	215	315	170
15	380	280	180	160	15	120	220	320	160
16	375	275	175	150	16	125	225	325	150
17	370	270	170	140	17	130	230	330	140
18	365	265	165	130	18	135	235	335	130
19	360	260	160	120	19	140	240	340	120
20	355	255	155	110	20	145	245	345	110

Notes: (i) The CE price is 250 and (ii) the CE volume is 40-41 units. (iii) The frequency of inframarginal (extramarginal) buyers' values and sellers' costs are 40-41 units (19-20 units) respectively. Thus, the frequency of inframarginal and extramarginal units are same across buyers' and sellers.'

In Market 1, the sum of the two highest value (lowest cost) units for each buyer (seller) are equal. In Market 2, the sum of values (costs) for buyer 1 (seller 1) is the highest (lowest) among all buyers (sellers). This sum of two highest value (lowest cost) units for buyers (sellers) goes down successively for the remaining buyers (sellers) such that it is lowest for buyer 20 (seller 20).

3.3 Competitive Equilibrium Predictions

A well-known definition of CE is the quantity Q* and price P* where the supply and demand schedules intersect (Marshall 1895, Walras 1877). In our markets, the intersection of demand and supply yields a CE prediction for the price at 250, and the predicted volume of transactions is 40-41 units. This CE price of 250 has implications for predicting trading activity and profits for individual traders. If the market price were always 250, it is in each buyer's (seller's) self-interest to trade their first 2 highest value (lowest cost) units (these units which should be traded because they earn a profit at the CE price are called inframarginal) and the third unit should not trade (and is called extramarginal), except for buyer 1 and seller 1 respectively. For buyer 1 (seller 1), the third unit is predicted to either be exchanged at a profit of 0 or not traded at all. Therefore, because the third unit of each trader either should not trade at P=250 or trade at zero profit, the value (cost) of the third unit is not predicted to have any impact on profit in either market. Given the differences in the distribution of buyers' values and sellers' costs across traders, differences in the distribution of profits among traders are predicted across markets 1 and 2. The CE prediction for profit for each trader is described in Tables 1 and 2. The CE prediction for distribution of profits across traders can be summarized as follows:

(i) Market 1: Complete equality in the distribution of profits (individual profits=205 per trader)
(ii) Market 2: Inequality in the distribution of profits (individual profits range from 110 to 300)

Note that Becker (1962) had shown that the "utility improving" choice behavior is sufficient to generate downward sloping demand and upward sloping supply functions in a market. Consequently, some experiments, have relied on these assumptions about the individual rationality of traders to simulate outcomes in DA markets. For example, Gode and Sunder (1993) assume that traders are budget constrained and therefore do not submit or accept offers that would result in losses. They describe this assumption as "market discipline," but it is a restriction imposed on traders such that their choices are always utility improving (rather than a stricter assumption like utility maximizing). This is consistent with Luce and Raiffa's interpretation of rationality as an attribute of all agents who attempt to take part in trades that do not diminish their utility (1957, pp. 192-193).

In the following section, we describe the strategies for robot traders that we shall utilize in our paper to simulate potential outcomes in DA markets.

3.4 Robot Traders and Strategies

Table 3 describes the treatments that we shall analyze to study the distributional outcomes. We shall analyze the market outcomes resulting from deploying the following strategies to traders:

1. Zero Intelligence traders

Each "zero intelligence" (henceforth ZI) trader generated random bids or offers (based on whether it is a buyer or seller) distributed independently, identically, and uniformly over a range of potential prices. Market discipline is imposed on trader behavior by restricting them from making direct (money) losses. Thus, ZI robot traders submit random bids and asks from a uniform distribution with support equal to its budget constraint. A buyer's bids are distributed $U[0, v]$, where v is buyer's value for the unit. A seller's ask is distributed over the seller's budget constraint $U[c, H]$ where c is the seller's cost for a unit and H is an upper limit of potential trading prices. With ZI sellers, a pre-determined upper limit to trading prices is necessary as a uniform distribution over the seller's actual budget constraint; $U[c, \infty)$ is ill-defined.

Table 3: Description of treatments

Trading Strategies	
ZI strategy	-A *ZI buyer* bids, i.e., makes an offer to buy, at a price randomly chosen from the interval $[0, v]$ where v is the buyer's unit redemption value set by the experimenter. -A *ZI seller* asks, i.e., makes an offer to sell, at a price randomly chosen from the interval $[c, H]$ where c is the seller's unit marginal cost set by the experimenter and H is a maximum permissible price in the market.
Kaplan Sniping	-A buyer waits until (current market ask ≤ redemption value) and either a low spread (current market ask − current market bid) ≤ target (set to 10) or a "mistake" (current market ask ≤ lowest price seen in the previous period), then send a bid equal to the current market ask. This causes a transaction. -A seller waits until (current market bid ≥ cost) and (current market ask − current market bid) ≤ target (set to 10) or (current market bid ≥ highest price seen in the previous period), then send an ask equal to the current bid. This causes a transaction.

Notes: (i) Kaplan Sniping has an additional "fail-safe" feature, which involves an additional strategy of trading near the end of the period by accepting the current market bid or ask if there is non-zero profit in doing so. (ii) Both ZI traders and snipers are budget constrained such that they do not accept offers that result in losses, i.e., strictly negative profits.

	Composition of Traders			
Configuration	*ZI traders (%)*	*Kaplan snipers (%)*	**Market Environments**	
1	100	0	Market 1	Market 2
2	95	5	Market 1	Market 2
3	90	10	Market 1	Market 2
4	80	20	Market 1	Market 2
5	70	30	Market 1	Market 2
6	60	40	Market 1	Market 2
7	50	50	Market 1	Market 2
8	40	60	Market 1	Market 2
9	30	70	Market 1	Market 2
10	20	80	Market 1	Market 2
11	10	90	Market 1	Market 2
12	5	95	Market 1	Market 2

Note: (i) The buyers (sellers) in each market are 20 in our experimental design. This yields a total of 40 traders in each market. (ii) At least one other buyer and seller type that submit bids and asks without relying on others, are essential to operationalize the Kaplan sniping strategy; which explains why we do not explore the market with 100% snipers.

Therefore, typically H is set at least as high as the highest buyer's value (Brewer et al. 2002).[18],[19] The use of ZI traders ensures minimal market discipline without any role for more sophisticated behavior.

2. Kaplan's snipers

The zero-intelligence trading strategy does not rely on sophisticated behavior that human participants could deploy in DA market. To address this deficiency, we also consider the impact of more sophisticated strategies on outcomes. Among other trading strategies, we consider "sniping" for the following reason: "Sniping" emerged the "winner" among a pool of 30 trading strategies (programs), in a high stakes tournament reported in Rust, Miller and Palmer (1993). This tournament was organized to isolate simple rule of thumb strategies that could be deployed by human participants in DA markets. They reported 30 programs that were submitted in their tournament: 15 were submitted by economists, 9 from computer scientists, 3 from mathematicians and remaining 3 from an investment broker, a professor of marketing and a joint entry from 2 cognitive scientists. Several of these programs emerged from working groups that co-developed sets of strategies. These groups include 7 entries from the Economics Science Lab (ESL) at University of Arizona, 3 from University of Minnesota and 2 each from University of Colorado (Economics) and Carnegie Mellon University (Computer Science) and 4 entries from the Santa Fe Institute. The best performing trading strategy (program), which we shall sometimes refer to as sniping, was submitted by the University of Minnesota economist Todd Kaplan that earned the highest profit out of 30 trading strategies reported in Rust, Miller and Palmer (1993). Given that sniping outperformed all other trading strategies that were submitted by human participants in Rust, Miller and Palmer (1993), we shall use sniping to capture sophisticated trading behavior in a DA market.

Here is a description of the sniping trading strategy: Kaplan robots follow a "steal the deal" strategy, waiting for a low bid-ask spread or a mistake before accepting profitable trades, and then executing a "failsafe" at the end of the period if there is a remaining opportunity. Thus, they require a market populated by other participants (other kinds of robots or humans) who actively submit bids and asks. In this trading strategy (Kaplan sniping): (i) buyers' wait until (current ask ≤ redemption value) and either a low spread (current ask − current bid) ≤ target or a "mistake" (current ask ≤ lowest price seen in the previous period), then send a bid equal to the current ask. This causes a transaction. Similarly, (ii) sellers' wait until (current bid ≥ cost) and (current ask − current bid) ≤ target or (current bid ≥ highest price seen in the previous period), then send an ask equal to the current bid. This causes a transaction. The fail-safe involves an additional strategy of trading near the end of the period by accepting the current bid or ask if there is non-negative profit in doing so. Operationally, we define the "end of the period" as ten

[18] For practical purposes, the support of the distribution from which the bids and offers could be made was restricted from (a) 1 up to the redemption value of the buyer and (b) seller's cost up to H for a seller.
[19] Note that because of budget restrictions, the distribution from which bids and asks are drawn depends on the induced values and costs. Therefore, the support of the distribution from which bids and asks were drawn randomly, is no longer identical across traders.

or less expected action opportunities for the trader to act under the Poisson arrivals distribution.

Thus, the sniping strategy (just like the ZI trading strategy) is non-adaptive, non-predictive, and non-optimizing and ensures budget discipline for each trader. Unlike the ZI trading strategy which is stochastic with distributional parameters depending on agent's values or costs, the sniping strategy is a pure non-stochastic function of agent and market conditions (including previous high and low prices and time remaining). The basic idea in the sniping program is to wait in the background while others do the negotiating, but when bid and ask get sufficiently close, or someone makes a big mistake (or offers a "great deal"), jump in and steal the deal.

Across market 1 and 2, the aggregate demand and supply conditions are fixed, but the variation in induced values and costs creates differences in CE predictions regarding the distribution of incomes across traders.[20] We test these predictions across markets 1 and 2, using a set of traders where each trader is assigned either of the two trading strategies: (a) sniping or (b) ZI trading strategies of Gode and Sunder (1993). Given that a variety of trading strategies could be deployed in the DA market environment, our approach to fixate on these two trading strategies has its limitations but, nevertheless allows exploring the boundary environments in which the distributional predictions of markets can be tested. This allows us to keep the underlying analysis tractable, helps identify scenarios where sniping does not outperform simple ZI trading strategy and make sense of the distributional outcomes which can be traced to varying levels of sophistication inherent in the trading strategies assigned to traders. [21]

Accordingly, we shall explore market environments, which are derived from varying composition of ZI and sniping traders. We successively infuse snipers in markets 1 and 2, from 0%, 5%, 10% up to 95% with the rest of the traders being ZI traders in any given market. Thus, markets, where there are 0% snipers, correspond to the benchmark environment with 100% ZI traders reported in Gode and Sunder (1993).

The second feature associated with infusion of snipers in markets is that we assign the sniper roles to traders who are likely to be most effective at sniping in successive configurations. This ensures that any treatment effect (due to the infusion of snipers) have the maximum potential effect on market outcomes and therefore any treatment differences are more likely to be observed. For example, a buyer sniper with value 100 is more likely to snipe successfully than a buyer sniper with value 20. Second, fixing the sniping strategy to specific traders allows us to track and examine the success of the sniping strategy with respect to the ZI trading strategy as compared to alternative ways in which sniping could have been exercised randomly by each trader. This is different

[20] We caution the reader against interpretation of induced values or costs as endowments. Induced values and costs yield profits only after successful trades at suitable prices and as such have no value if the corresponding units are not traded. It would perhaps be more appropriate to think of induced values and costs as quasi-endowments or endowments of options to sell or buy units rather than endowments of units.

[21] We hope that the analysis developed in the paper will ultimately be exploited to analyse outcomes that may arise from other potential trading strategies and any combination of trading strategies in each environment.

from an alternative approach where traders could have been programmed to snipe with a certain probability in each instance and perhaps could have used some adaptive rule to switch between sniping and ZI trading in response to observed outcomes. Such alternatives go beyond our desire to work with simple, non-adaptive, non-stochastic and non-optimizing strategies as a first cut to investigate the main questions.

Third, since the strategies we explore here are non-adaptive, the reader is cautioned against any interpretation of our results as those derived from a dynamic process. It would perhaps be more helpful to think of the different period outcomes as different instances of operation of the underlying markets, even though there is some previous period dependence introduced through the trading rules by which snipers operate. As we shall discuss later, it is more appropriate to test the distributional predictions by aggregating the profits from all these instances. However, the aggregation of profits should not be interpreted as those derived from some dynamic historical process.

4. Results

Given the stochastic nature of market trading, that can be attributed to the random strategies of ZI robot traders and the random arrivals of individuals there is significant variation in prices, overall volume of transactions and accumulated profits from period to period. We shall report the results from 10,000 periods for each market treatment.[22]

As described above, we successively assign the role of snipers to buyers and sellers in the order of id number. This is the same ordering as the top-down ordering of Table 2 describing cost and values in markets 1 and 2. Thus, in the market with 5% snipers, buyer 1 and seller 1 are snipers; in markets with 10% snipers, buyers-1 and 2, and sellers-1 and 2 are snipers, and so on.

Our assignment of sniping strategy implies that buyers and sellers that have the highest ex-ante CE profit levels in Market 2 deploy the sniping strategy as the volume[23] of snipers in the market is increased successively.[24]

Example Trading Periods

Figures 2 and 3, below, present a single trading period example for a market (Market 1) populated entirely by ZI traders (Fig. 2), and a market populated by 19 sniper buyers and 19 sellers and the remaining agents, 1 buyer and 1 seller, are ZI traders (Fig. 3). Each time series displays experiment time, on the horizontal axis, and price on the vertical axis. Each single period example has a single period beginning at t=1000 and concluding just before t=2000. Bids and asks are colored dots, and the series of transaction prices is connected by a line. Recall that the price of a transaction is always determined by an earlier order, so it is common to see high bids triggering a transaction at a lower price, from the earlier pre-existing best ask, and it is common to see low asks triggering a transaction at a higher price, from the earlier pre-existing best bid.

[22] This will become clear when the distribution of Gini coefficients is discussed for various treatments.
[23] Clarification: The *volume of snipers* in the market refers simply the number of Kaplan Sniper agents and not the number of trades by these agents.
[24] This is trivially true in market 1 in which all traders have equal ex ante profit in the competitive equilibrium.

In the 100% ZI traders market, a typical single period consists of dozens of trades (recall, CE predicts 40-41 transactions). In the example, most of these occur early in the period (only one trade after t=1200), and it is difficult to separate the individual trades visually in a figure that shows the entire period. The trade prices over time follow typical ZI market behaviour. The high initial price variance occurs because buyers with high values or sellers with low cost are more easily matched for a trade very early at some price, and it could be a price far from CE. The convergence over time towards CE occurs because the high profit trades are eliminated early, and the low profit trades occur later, involving agents with values and costs closer to CE, and tend to require numerous random generations of bids and asks to yield an overlap resulting in a trade.

In the sniper dominated market 1, a single period has a much smaller number of trades, and we see the failsafe trades that Kaplan snipers make at the end of the period occurring after t=1900. Recall that since the Kaplan snipers always send bids and asks that accept the pre-existing bid or ask in a trade, the only trades that can be made by a Kaplan Sniper are with the single ZI trader on each side of the market. This means there can be at most 6 transactions between the ZI traders and Kaplan Sniper traders, and fewer than 6 transactions if the two ZI traders trade with each other. The last two trades in Figure 3 happen near the end of the period and at very different prices. As will be discussed later, this is a typical market consequence of the Kaplan Sniper fallback strategy for snipers who were unable to find a unit at advantageous prices or with low bid-ask spread during the first 90% of the time in a period.

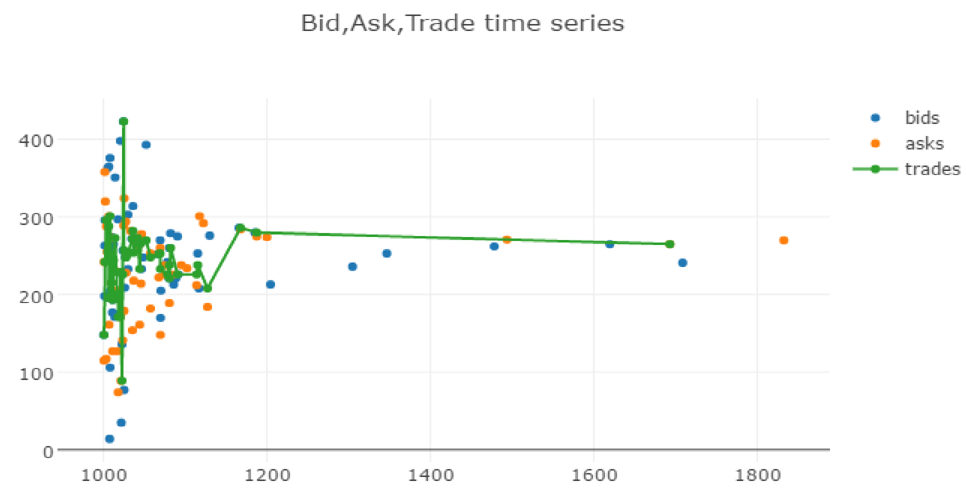

Figure 2: Example trading period- 100% ZI traders – Market 1

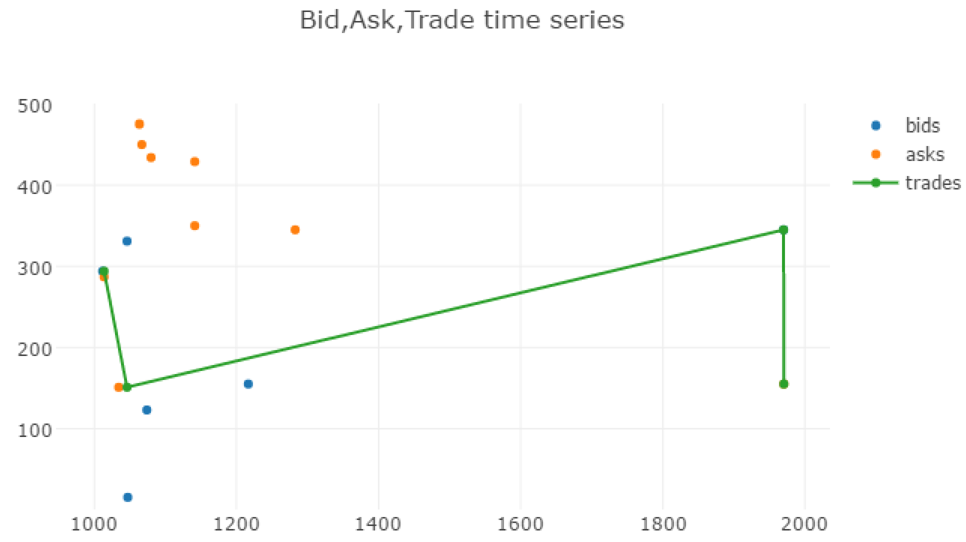

Figure 3: Example trading period- 95% Kaplan Sniper, 5% ZI – Market 1

Next, we shall discuss the observed outcomes concerning prices, allocative efficiency and distribution of profits.

Prices

Figure 4 describes the box-whisker plots for the distribution of observed prices for the 12 different configurations, which correspond with various volumes of snipers in market 1 and market 2 respectively.[25] Note that the configuration with 0% snipers is the benchmark environment, which corresponds to the market environment with 100% ZI traders reported in Gode and Sunder (1993). The median observed price for market 1 and market 2 is close to the CE predicted price of 250 for all configurations. There is virtually no difference in the median price for the two markets across various market configurations.

The average (mean) observed price for each market and configuration (volume of snipers in the market) are shown in Figure 5. Note that the mean observed price for markets 1 and 2 all lie between the range 249-251 which includes the CE predicted price of 250. Visually, there does not seem to be any systematic difference in mean observed price for markets-1 and 2-as the overall volume of snipers varies across configurations. This is consistent with the differences in the median observed prices reported in Figure 4.

[25] The confidence intervals plotted in the box plots in all figures in the paper are the 50% confidence intervals with the median shown as the solid line inside the box bars and values outside the bar were ignored for visual clarity.

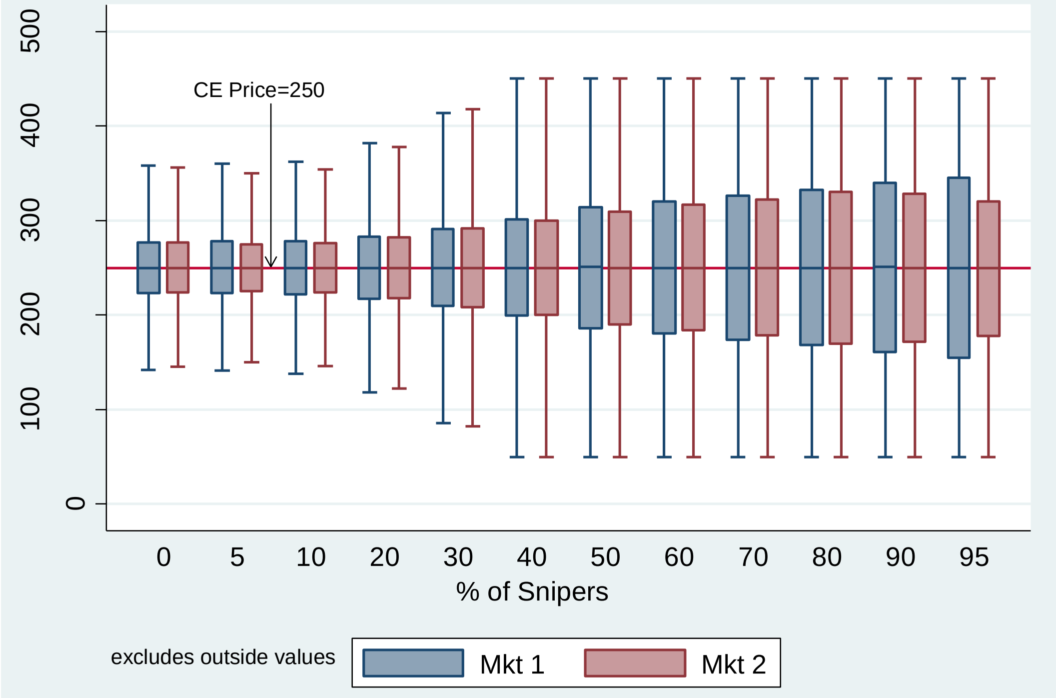

Fig 4: Distribution of Observed Price by market

To test whether these differences are significant, we report regression estimates for observed prices on a market dummy (“*mkt2*”=*1* for market 2; 0 otherwise), dummies for configurations (“$config = x$” $= 1$ for configuration “x”; 0 otherwise) corresponding to various volumes of snipers and their interactions. These results are reported in column 1 of Table 4.

Note that the estimated coefficient for the dummy for *mkt2* is positive but insignificant at 10% level, which indicates that in a market with 100% ZI traders, the mean observed price in market 2 is slightly higher than in market 1. The coefficients attached to most configuration dummies are not significant except for $config = x$ $(x = 5, 30)$. This suggests that in market 1, the differences in the overall volume of snipers do not have a significant impact on mean prices. The coefficients for the dummies for interaction terms “$mkt2 \times config = x$” $(x = 5 \; and \; 10)$ are negative and significant at various levels.

The coefficient for the constant term is not statistically different from the CE prediction of 250 ($p \; value$ for $F(1, 7725027)$ equals 0.5743). This yields the following result:

Result 1: *(i) The observed average prices across markets-1 and 2- are distributed around the CE price of 250 for various configurations of snipers infused in the markets. (ii) There are no systematic differences in observed average prices in markets 1 and 2 as the volume of snipers varies across markets.*

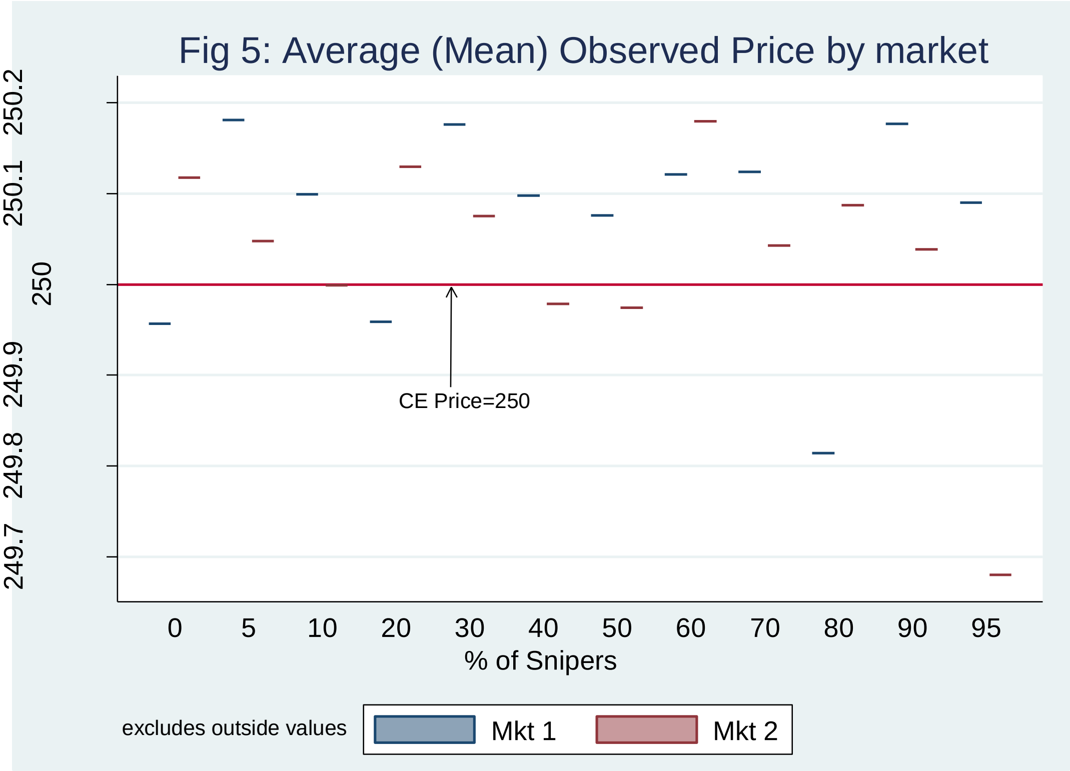

Table 4: Regression estimates for treatment differences

	(1) Price	(1) Efficiency	(1) Gini
mkt2	0.16	0.00	0.04***
	(0.11)	(0.01)	(0.00)
config=5	0.22**	-0.22***	-0.00***
	(0.11)	(0.01)	(0.00)
config=10	0.14	-0.77***	-0.00***
	(0.11)	(0.01)	(0.00)
config=20	0.00	-3.30***	-0.00*
	(0.11)	(0.01)	(0.00)
config=30	0.22*	-7.95***	0.02***
	(0.12)	(0.02)	(0.00)
config=40	0.14	-14.47***	0.08***
	(0.13)	(0.02)	(0.00)
config=50	0.12	-22.50***	0.16***
	(0.15)	(0.03)	(0.00)
config=60	0.16	-37.66***	0.30***
	(0.17)	(0.04)	(0.00)
config=70	0.17	-53.86***	0.44***
	(0.20)	(0.03)	(0.00)
config=80	-0.14	-69.45***	0.56***
	(0.23)	(0.03)	(0.00)
config=90	0.22	-84.40***	0.69***

	(0.31)	(0.02)	(0.00)
config=95	0.13	-91.61***	0.75***
	(0.41)	(0.02)	(0.00)
mkt2 # config=5	-0.29*	0.03**	0.00**
	(0.15)	(0.01)	(0.00)
mkt2 # config=10	-0.26*	-0.21***	-0.00***
	(0.15)	(0.02)	(0.00)
mkt2 # config=20	0.01	-2.02***	-0.01***
	(0.15)	(0.02)	(0.00)
mkt2 # config=30	-0.26	-4.58***	-0.02***
	(0.16)	(0.03)	(0.00)
mkt2 # config=40	-0.28	-7.02***	-0.04***
	(0.18)	(0.03)	(0.00)
mkt2 # config=50	-0.26	-8.30***	-0.06***
	(0.21)	(0.04)	(0.00)
mkt2 # config=60	-0.10	-6.07***	-0.09***
	(0.24)	(0.05)	(0.00)
mkt2 # config=70	-0.24	-4.61***	-0.10***
	(0.27)	(0.05)	(0.00)
mkt2 # config=80	0.11	-3.01***	-0.09***
	(0.32)	(0.04)	(0.00)
mkt2 # config=90	-0.30	-1.15***	-0.08***
	(0.43)	(0.04)	(0.00)
mkt2 # config=95	-0.57	-0.15***	-0.07***
	(0.56)	(0.03)	(0.00)
Constant	249.96***	98.99***	0.20***
	(0.08)	(0.01)	(0.00)
Observations	7725027	240000	240000
R^2	0.0000	0.995	0.985

Notes: (i) Robust Standard errors in parentheses (ii)* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Efficiency

To investigate allocative efficiency across markets, we calculated average efficiency per period, where efficiency in a market period -- as first described in the experimental literature in Plott and Smith (1978) -- is the ratio between total realized profit from the empirical data and maximum possible total profit. This maximum possible total profit is the same as the total of theoretical consumer and producer surplus at CE[26].

$Efficiency = Total\ Profit\ Realized/Maximum\ Profit\ attainable\ in\ CE$.[27]

Figure 6 describes the box-whisker plots for the distribution of observed average efficiency for the 12 different configurations, which correspond with various volumes of snipers that make up those markets.

[26] For this explanation see the review article in Plott(1982), pp.1491-1492.
[27] Note that, in these markets, the maximum attainable profit irrespective of agent strategy is the profit predicted by CE.

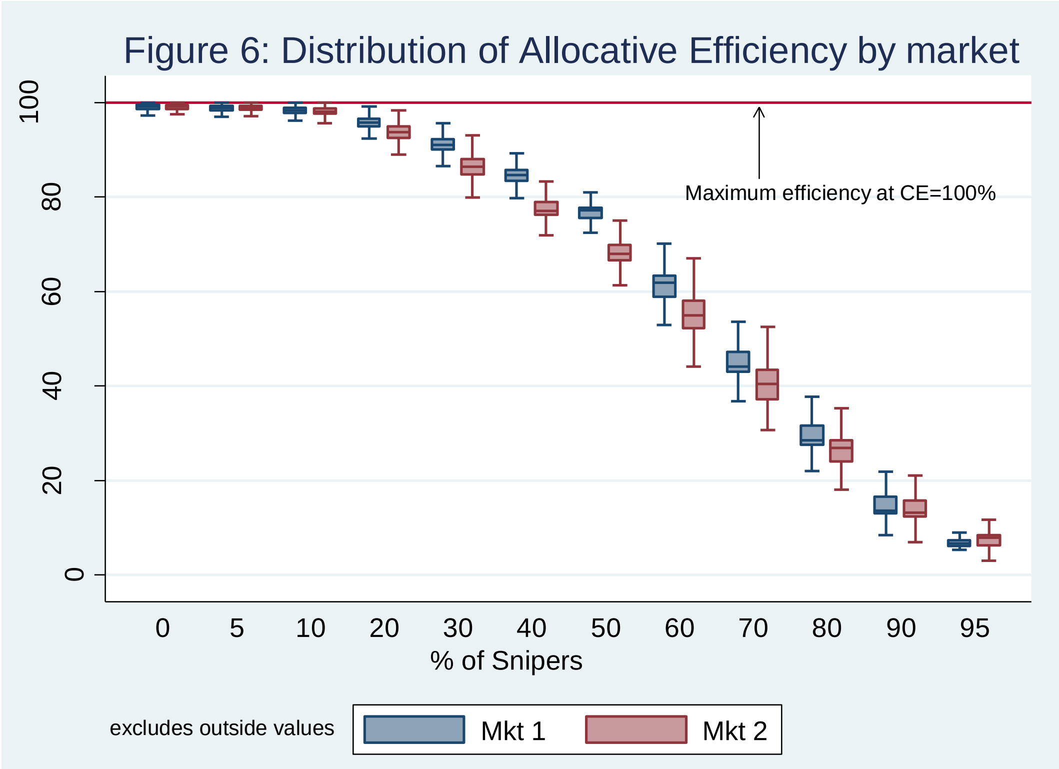

As evident from Fig 6, the average (median) efficiency is close to 100% in markets where the volume of snipers in the markets are 10% or less. In markets where the volume of snipers is 20% or more, the average (median) efficiency is less than the maximum possible levels of efficiency. The mean efficiency levels for these markets are described in Fig 7. The average mean efficiency levels in market 2 are similar to those in market 1; but as the volume of snipers in the markets goes beyond 10%, the mean efficiency level is higher in market 1 than market 2.

To test whether these differences are significant, we report regression estimates for efficiency on a set of explanatory dummy variables. These results for average efficiency are reported in column 2 of Table 4. The estimate of the dummy for market 2 is close to zero and insignificant. The estimates for various configuration dummies "$config = x$" ($x = 5, 10, 20, 30, 40, 50, 60, 70, 80, 90$ and 95) are negative and highly significant. This indicates that the observed average efficiency becomes successively smaller as more snipers are infused in the market.

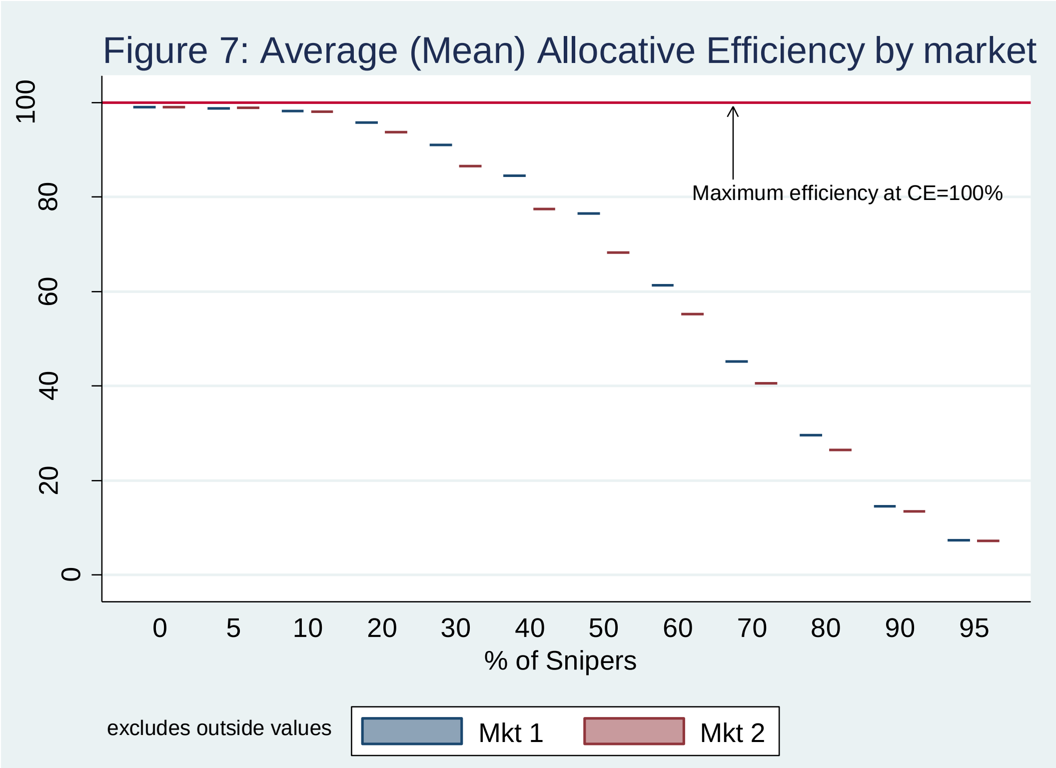

The estimates for the interaction dummies for "$mkt2 \# config = x$" are negative and highly significant for $x = 10, 20, 30, 40, 50, 60, 70, 80$ *and* 90; and this estimate is positive for $x = 5$ and significant at 5% level. This gives us the following result.

Result 2: *(i) The observed allocative efficiency is very high and close to the CE prediction in markets where the volume of snipers is 10% or less. (ii) The observed mean allocative efficiency in markets-1 and 2-decreases, as the volume of snipers in these markets, increases with slightly higher efficiency in market 1 with respect to market 2.*

Next, we shall discuss the distributional outcomes in the two markets.

Observed Inequality and distribution of Profits

Our design allows us to track the profits accrued to each trader. Our approach to analyzing the distributional outcomes is different from the approach where the distributional outcomes are analyzed from the perspective of buyers or sellers as groups (e.g., see Smith and Williams 1982). We take an approach that is more often utilized in the development literature, where distributional outcomes are analyzed in terms of Lorenz curves which captures the cumulative shares of profits across individuals; this is often used to derive a measure of inequality called the Gini coefficient. The Gini coefficient is defined as the ratio of the area between the Lorenz curve and the 45-degree line to the area under the 45-degree line, which indicates complete equality of distributional outcomes. A value of zero for the Gini coefficient corresponds to complete equality whereas a high value (approaching 1) of the Gini coefficient corresponds to a high concentration of wealth and therefore high levels of inequality.

CE predictions for Inequality (Gini Coefficient)

By construction, the CE predictions for Gini coefficient follow:

(i) The predicted profits for all traders are equal in market 1. This implies that the CE prediction for Gini coefficient in market 1 is zero; and

(ii) There is some variation in the predicted profits for traders in market 2. Therefore, the CE prediction for Gini in market 2 is positive and equals 0.1664.

First, note that it is possible to derive the Gini coefficient in each period from the distribution of profit across all traders in that period. Figure 8 shows the Box-Whisker plots of Gini coefficients for various configurations of snipers in markets-1 and 2-for the 10,000 periods. The red horizontal line denotes the CE prediction for Gini in each market. However, given that there is significant variation in distributional outcomes across periods due to the randomness of the underlying trading activity, the CE predictions for Gini must be compared with Gini that is derived from the profits accrued by traders over 10,000 periods. We enhanced the box-whisker plots in Figure 8 by adding the Gini coefficient derived from accumulated profits in 10,000 periods, which is shown as a long dash. Note that, the Gini based on accumulated profits absorbs the random variation in profits that is observed in individual periods; therefore, the Gini based on accumulated profits is much lower than the Gini in a typical period.

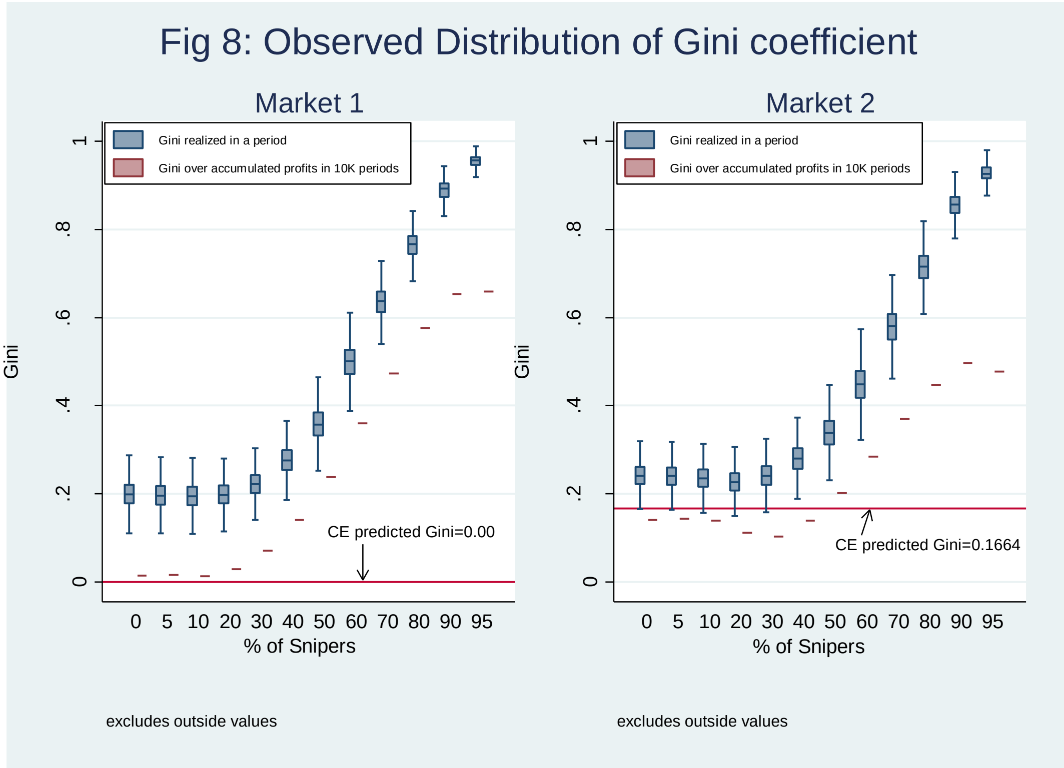

The left and right panels of Figure 8 display the distribution of Gini coefficients for market 1 and 2 respectively. Figure 8 suggests the following:

(i) The Gini derived from accumulated profits in 10,000 periods is closer to the CE predictions for markets with a low volume of snipers and the Gini derived from accumulated profits is much higher than the Gini predicted by the CE for markets with a higher volume of snipers (50% or more snipers).

(ii) The Gini derived from accumulated profits is lower than the CE predictions in market 2 with a low volume of snipers (40% or fewer snipers).

To further investigate the differences in Gini coefficients in these markets, we regress the observed Gini in a given period on explanatory dummy variables for -market type, the configuration of snipers and their interaction terms. The results for this regression are reported in column (3) (Table 4). The estimate for the constant term is positive and highly significant ($p\ value < 0.01$). The estimate of the coefficient for *mkt2* is positive and highly significant. The estimate for the coefficients for $config = x$ dummies are positive and highly significant (except for $config = 5, 10\ and\ 20$). These results indicate that, with 100% ZI traders, the mean Gini is higher for market 2 with respect to that for market 1. The estimates for the coefficient for the interaction dummies are mostly negative and highly significant (except for $mkt2 \times config = x; x = 5$). This suggests that as more snipers are introduced in the markets, the mean Gini is lower for market 2 with respect to that for market 1. By comparing the Gini derived from accumulated profits from 10,000 periods in each market, the following result is obtained.

Result 3: *(i) The Gini coefficient derived from accumulated profits in each market is roughly similar to the CE predictions for each market at low* volumes *of snipers in these markets. The observed inequality based on Gini, derived from accumulated profits, for market 2 is lower than the CE predictions at low volume of snipers (40% or* fewer *snipers); (ii) The Gini coefficient derived from accumulated profits in each market is higher than the CE predictions for each market at high levels of sniper infusion in markets (50% or more snipers).*

Distribution of profits: Snipers versus ZI traders

We can further investigate the relative performance of snipers versus ZI traders to delineate the environments in which one type of traders outperforms the other types.

To understand the relative performance of snipers versus ZI traders, we define a variable “Ratio1” which is the ratio of the mean accumulated profit in 10,000 periods divided by the profit predicted in the CE for each trader (taking into consideration all three units for each trader). If this ratio is above 1, it indicates that the mean profit accumulated by a trader is higher than that predicted under the CE. If it is below 1, it indicates otherwise. A value of 1 for “Ratio1” indicates that the mean accumulated profit is equal to the CE prediction.

Figure 9 displays the box-whisker plots for the variable “Ratio1” for the various types of traders across various market configurations. From this figure, it can be inferred that the snipers perform better with respect to the CE predictions in scenarios where the volume of snipers in the markets is low rather than high in each market. In market 1, since the expected profit under CE is equal for each trader, the distribution of “Ratio1” allows us to compare the performance of snipers versus ZI traders. In the left panel

(Figure 9), the box bars for snipers shrink and become smaller in vertical height, as the volume of snipers is 20% and above. This implies that snipers perform worse than (i) the CE predictions and (ii) ZI traders (in market 1), as the volume of snipers is 20% and above.

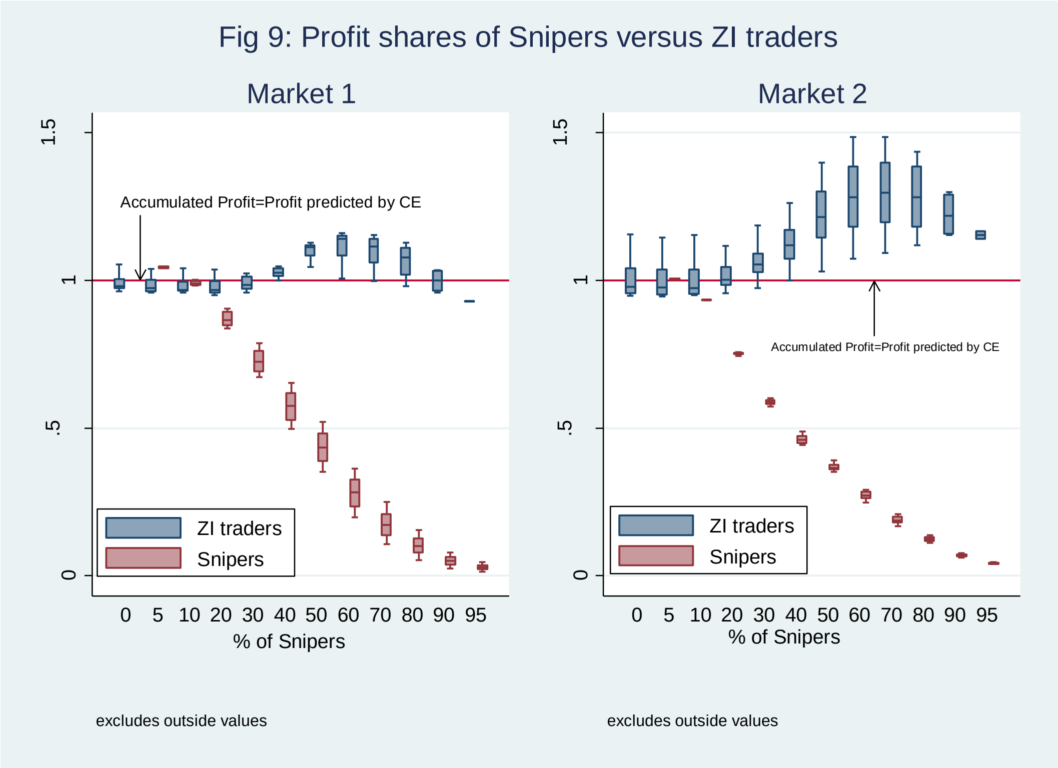

The box-whisker bars for ZI traders shift upwards and the medians observed are larger than 1 as the volume of snipers is 20% and above. This suggests that, in markets where the volume of snipers is 20% or more, the ZI traders perform better than that predicted by the CE across both markets and their performance is better as compared to snipers in market 1. Since the ZI strategy itself is not changing, this increased performance occurs from the ZI agent's trade with more Kaplan Sniper agents who are trading the "fallback strategy" at the end of each period.

In market 2, a direct comparison between snipers and ZI traders is not possible, since the expected profits in a CE vary across traders. However, it can be inferred that snipers perform much worse with respect to the CE predictions than the ZI traders as the box bars for snipers shrink and shift downwards as the volume of snipers is 20% and above. To compare the performance of snipers versus ZI traders in market 2, we define another variable. Note that in our design, there is no difference between seller 1 and buyer 1 based on the predicted CE profit for these traders. This allows us to normalize the accumulated profits for each trader by the sum of induced values as we go down the list in Table 2 while treating seller 1 as buyer 1 and so on. The variable "Ratio 2" is the ratio of accumulated profits to the sum of induced values is derived for each trader. This is different from the variable "Ratio1" shown in figure 8 where the accumulated profits are compared to CE predicted profit and not induced values. Figure 10 displays the box-whisker plots for the variable "Ratio 2" in market 2 for all traders.

From Figure 10, it can be inferred that the performance of snipers versus ZI traders worsens as the volume of snipers is 30% and above.

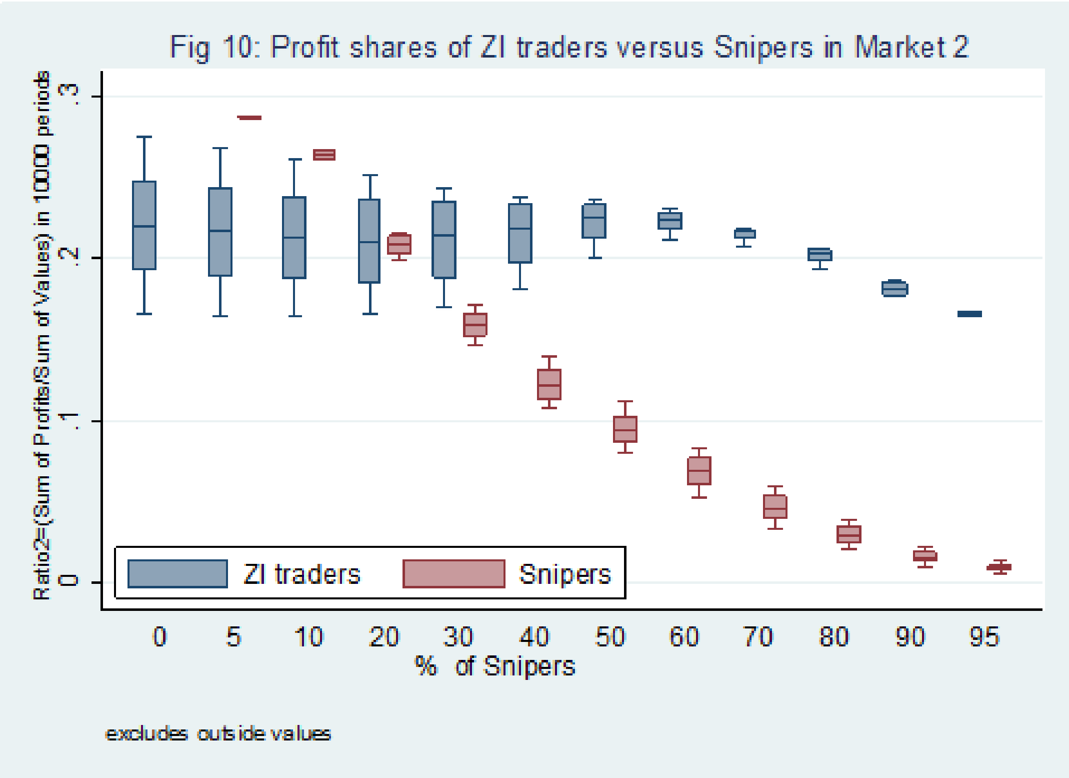

On average, snipers perform poorly with respect to ZI traders in markets where the volume of snipers in the market is relatively high. This deserves a question regarding the deployment of a sniping strategy in markets where the volume of snipers is relatively high. To select from a set of alternative strategies to maximize profits requires more information at the level of a trader which may not be accessible in a DA market where each trader knows only their induced values before trading occurs. In other words, the profitability of sniping strategy to various other alternative strategies cannot be easily inferred beforehand by an individual trader.

In a DA market, traders might select from alternative strategies based on the relative profitability of successful trades, in the absence of more information about the private outcomes for other traders to compare the profitability of alternative trading strategies in a dynamic environment. Next, we show that the profit per successfully traded unit (as a fraction of induced value) is higher for snipers than ZI traders as the volume of snipers rise in each market. As above, we treat seller 1 as buyer 1 and so on to normalize the profits from successful trades. We define a variable "Ratio3" which is the ratio of profit earned on a successfully traded unit to the induced value for that unit. Figure 11 displays the box-whisker plots for snipers versus ZI traders for the variable "Ratio3" for each market. The box bars for snipers shrink and shift downward as the volume of snipers in the market rises. In market 1, the box plots suggest that the profitability based on successful trades for snipers is higher than ZI traders until the volume of snipers reaches 50%. In market 2, the profitability based on successful trades for snipers stays above the profitability for ZI traders for most configurations, which correspond with various volumes of snipers in the markets.

Fig 11: Profit shares of Snipers versus ZI traders on successful trades

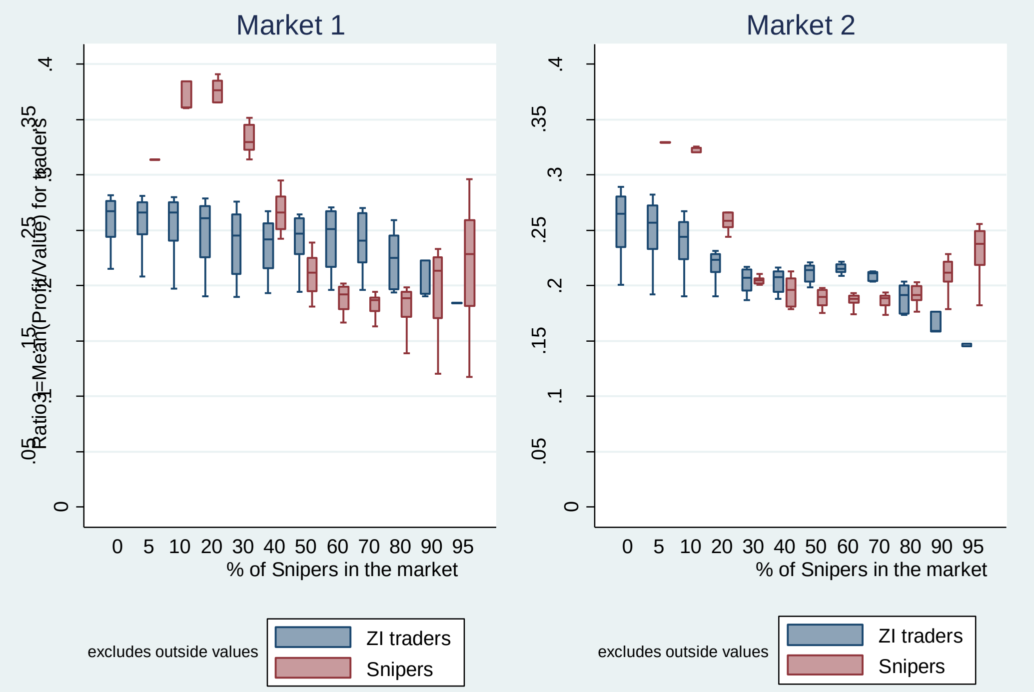

This suggests that in markets where the volume of snipers in the market is relatively high if snipers do not consider the profits foregone on unsuccessful trades and in the absence of information about outcomes for other traders, switching to alternative strategies may not be obvious.

Further analysis of what is happening in the Kaplan Sniper dominated markets is hinted at by the divergent end-of-period prices shown in the trade time series of Figure 3. Near the end of a period, all of the trades possible among the ZIs have probably occurred, and all the low-cost/high-value units among the ZI exhausted, leaving the snipers fallback strategy to be satisfied by trading for the ZIs extramarginal units. Further analysis is possible from a careful examination of tick-by-tick market behavior, but at some point, distracts rather than adds to the examination of the broad general questions about markets and inequality described in section 1.

5. Reflections

We now return to the general questions posed in Section 1 that motivated our investigation and consider whether the results presented in the previous section can help shed light on these broad issues. We do not expect any of these issues to be rigorously solved by any single study. For some of these general questions there is progress; but in others, the answer may be suggested by refinements or future work.

1. Do market processes worsen inequality or merely duplicate inequality already inherent in the economic environment?

We find that in this study, in the long-run market processes alone do not worsen inequality beyond the predictions of prices and allocations from neo-classical economic theory -- but rather there are short-term inequality-worsening effects from (i) randomness in trading conditions, and (ii) harmful trading strategies. The Gini-coefficient analysis of Result 3 / Figure 8 suggests that in the long-run, or aggregate, that the DA market with ZI traders tends to duplicate the inequality predicted by the theoretical competitive equilibrium (CE) allocation. Inequality, as measured by Gini coefficient, is higher when (a) considering individual periods of the market in isolation, and (b) introducing a "harmful" trading strategy by replacing ZI traders with the Sniping strategy.

While we can speculate that adding features like a periodic "cost of living" to the agents' environment together with high-interest loans for covering shortfalls would generate higher Gini coefficients and more "unfair" outcomes, adding these properties of field scenarios is not necessary to generate differences in short-run/long-run inequality. Indeed, adding such features too early in modelling, as we have avoided, could conflate multiple phenomena.

2. Does the neoclassical theory … succeed to explain the distribution of incomes?

Yes, but only in the markets with mostly ZI traders. Although from Fig.1 and Result 1 we show that the average price in all markets, regardless of ZI vs. Sniper populations, is consistent with the CE price, which if constant would lead to the distribution of incomes consistent with the distribution of incomes predicted by the CE. However, there is considerable variance in price time series within a period. This is evident from the size of the price distribution bars in the boxplot of Figure 4 and causes the divergence of individual period Gini coefficients from those calculated over aggregate profit from all periods. Also, the snipers cause a deviation from a simple supply/demand model as follows: Failure of snipers to find a "good deal" causes them to execute their fallback strategy of trading near the end of the period, at whatever prices are then available from the remaining ZI traders. When there are many snipers, this rush to trade near the end of the period and inability to trade with other snipers results in a net transfer of income from the snipers to the remaining ZI strategy traders that can be inferred from Figure 9 and the end of the time series of Figure 3. This happens to both the buyers and sellers in the market and does not move the average price but rather causes deviations of prices away from the average and in both directions depending on who is buying and selling. It is possible that neoclassical theory could explain this income transfer within a suitable reparameterization. Such a project is left for future work.

3. Can we learn something about the inequality in the real world by observing the outcomes generated by market processes that do not discriminate between agents except for their cost and value endowments? Is the inequality observed in the real world primarily produced by institutional features or market processes?

The latter part of the question often is built on a kind of a critique or pre-judgment that assumes the experimental environment is somehow not real even though it is a subset of the real world. Previous experimental economists have also grappled with this issue and can be found in methodological and conceptual explanations of Charles Plott, Vernon Smith and others. If there are economic laws that are generally applicable, they apply to both to experiments and to complex historical and current events. We mention in the introduction that because of the complex milieu surrounding world and national economies, additional insights can be obtained in a more simplified environment. Here we have constructed one such environment. We were able to generate inequality without confounding factors such as discrimination in capital and labor markets, unequal opportunities in education, effects of government intervention, or environmental factors.

The market process operating in this environment, that is, the rules by which bids and asks are processed into trades, does not discriminate between traders. Thus, it can be inferred that trader properties and behavior, not the behavior of the market, drives the long-term level of inequality in our environments. This includes not only the differences in trader assignments of costs and values, to which the CE model is sensitive but also aspects of agent trading strategy captured by replacing some of the ZI traders with Sniper traders.

We explained earlier that with regard to trader assignments of costs and values, the aggregate values and costs in Market 1 (equal profit in CE) and Market 2 (unequal profit in CE) are identical. The CE prices and quantities are identical. It is the order of assignment of unit costs and values among the agents that produce differences in theoretical profits that is later observed in action in the data. In particular, in Market 2, the agents' profit is highest for low id numbers. What may be novel is not merely that the empirical data from the robot markets confirms the theory, but it demonstrates that a change in inequality among agents does not need to involve a change in prices. Also, the change in inequality as measured by a Gini coefficient is sensitive to whether we are talking about the long run or short run, and the Gini coefficient may hide complexities that can only be uncovered by more detailed analysis (such as why the Gini coefficient can initially decrease as we add snipers).

The Sniper traders' strategy results in discrimination, but against similar snipers. This is a different kind of discrimination from "real world" discrimination whereby members of a group boycott or refuse to trade with others who are different in race, religion, or some other quality. One could say the snipers also refuse to trade fairly with the ZI traders, given that snipers have adopted a strategy that is designed to primarily pursue opportunities that are either likely trading "mistakes" (bids or asks that are too generous compared to CE) by the randomly-behaving ZI agents, or rare extreme price fluctuations, or only low-risk with a narrow spread, all while leaking as little pricing information as possible. Similar strategic elements are also seen in two other common cases: (i) in labor markets where a potential employer (a labor buyer) provides no salary estimate but expects acceptable candidates (labor sellers) to announce the first number in terms of a salary request, hoping to find a candidate who does not know the market value of their work and asks too little; and (ii) on stock and commodity exchanges as immediate market orders that remove liquidity from the order book, where they are often discouraged through higher fees vis-a-vis queued limit orders that add liquidity.

Failures of the snipers in our markets to obtain a superior outcome through speculation on the possibility of "great deals" at prices above (below) the previous period high (low) often causes them to lose income to the ZI traders in the fallback strategy of trading near the end of the period. Sufficient numbers of snipers decrease market efficiency to the degree that it harms all traders. We think Pope Francis was probably referring to bets on the direction of future price movements in financial and commodity markets when he spoke of "speculation," but the Sniper traders in this study also show an example of market speculation gone awry. The snipers demonstrate the potential harm of placing undue weight on the likelihood of favorable future events in pursuing self-interest, and it is noteworthy that the sniping strategy harms traders as a whole and not only the snipers.

Finally, random aspects of the market process, such as varying the sequence in which traders place their bids and asks, appears to be part of the explanation for why short-term inequality is higher than long-term inequality. In future work, other forms of markets, such as uniform price call markets, could be tested where the simultaneous determination of price for batches of trades could further reduce price variance.

As to the inequality proportions part of question 3, while we obviously cannot at this stage determine what proportion of world inequality can be attributed to institutional versus market processes, the fact that we can reproduce some inequality with simple processes is an important step towards investigating more complex scenarios in the future.

4. Is there some natural pattern that emerges in a macro sense but not in a micro sense (I.e., we all experience good and bad times, but there is no inherent bias)? Do markets systematically amplify the fortunes of the rich and hurt the poor? By contrast, do markets protect people?

We do find evidence of such a macro pattern, such that inequality "averages out" in the long run towards CE allocations with the ZI traders. Equal allocations in CE yield equal long-run average allocations. This is true even though the individual periods have varying allocations, i.e., "good times and bad times" for specific individuals, that average out. Market trading alone does not systematically amplify the good or bad over time. We do not have capital markets in our model, and some capital market practices (e.g. "payday loans" and "loan sharks" at high-interest rates, targeted at the poor) together with a "cost of living" could be introduced into these models to possibly result in divergent aggregate outcomes. Speed may be another variable of interest with regard to the snipers' strategies that could create an additional bias. We have not given the snipers a speed advantage, as is sometimes a concern in finance (e.g., with high-frequency trading in the stock markets). This is left for future research.

Similar to the results in Gode and Sunder (1993), we find that markets can protect agents. The replacement of the ZI traders by snipers shows that "unfair" strategies like sniping can fail to generate extra profits for the snipers. The ZI traders are protected from the snipers' plan to execute trades at unfavorable prices in part by market competition from other ZI traders. Additionally, even the snipers are protected from earning zero profits from the failure to find a "great deal" by the possibility of trading at available prices in the market near the end of the market period.

6. Conclusions

The paper takes an important conceptual step towards understanding the distributional consequences of DA markets while controlling for other influences. It is important to disentangle the relative contributions of markets and other institutions to the levels of inequality that have been documented in recent times to improve our understanding of markets and inequality.

The existing literature has shown that in markets where traders are endowed with minimal intelligence, convergence to a unique CE price and very high levels of allocative efficiency are observed. The distributional questions have been previously analyzed to the extent that the ratio of the producer surplus to the consumer surplus determines whether prices converge to the CE price from above or below. The convergence of prices to CE predictions, and especially when averaged or aggregated to reduce variance, however, can mask the distributional consequences of market processes. We go beyond the literature by taking two steps: first, we analyze environments other than those in which all traders are endowed with minimal intelligence and second, we analyze how the distribution of profits compares with the CE predictions of inequality. We analyze environments in which the traders are either ZI or snipers, and their relative volumes vary in the markets. The environments in which snipers are gradually infused while the volume of ZI traders shrinks may be viewed as “boundary experiments” (Kaplan 1964) similar to those pursued in Smith and Williams (1990).[28]

The predictions of the CE are consistent with observed outcomes for aggregate prices, efficiency levels and the Gini coefficient for accumulated profits in markets where the infusion of snipers is relatively small (10% or less of the overall volume) with respect to the ZI traders. This includes the benchmark environment with 0% snipers (or 100% ZI traders) in markets. In market 1, which is designed to produce complete equality in the distribution of profits, the DA market approximates that outcome. In market 2, where some ex-post inequality in the distribution of profits is chosen, the realized distribution of profits is still consistent with the predictions of the CE. In markets, where the volume of snipers is sufficiently large (20% and above of the overall volume), the observed allocative efficiency declines with respect to the CE prediction and the Gini coefficient based on accumulated profits indicates a much higher level of inequality than predicted by the CE. These outcomes are observed for both markets 1 and 2.

Based on the environment with 100% ZI traders, Gode and Sunder (1993) concluded that consistency of the outcomes with the CE predictions is primarily driven by the market institution and the “the effect of human motivations and cognitive abilities has a second-order magnitude at best.” Our results suggest that this claim does not extend to environments where the volume of snipers goes beyond a certain threshold.

[28] “*Boundary experiments are explicitly associated with some set of laws and consist of fact-finding enquiries designed to fix the range of the application of the laws, particularly with regard to extreme conditions. In behavioral science, such experiments are exemplified by studies of sleeplessness, sensory deprivation, perceptual thresholds, and the like.*” (Kaplan 1964).

The impact of trading strategies on market opportunities for other traders and the volume of successful trades needs to be discussed to reflect upon the distributional outcomes. The ZI strategy, while non-strategic, through its random trading behavior contributes a kind of beneficial speculation that lubricates trading by offering a price.

In contrast, the snipers capitalize on mistakes made by others, and in environments where their volume is low compared to the rest of the market, outperform the ZI traders. However, when the number of snipers rises, their strategy begins to impede beneficial trades; this shows up in the lower allocative efficiency levels in these markets.

The snipers sometimes do better on individual trades, but less frequently as the number of snipers increase, and also, they trade less frequently overall. It is easy to imagine that in extreme cases, for example, if very few traders make mistakes early in the trading period or if there are too many snipers in the market, the outcomes observed are different from those observed markets with 100% ZI traders. In those circumstances, sniping, which is a more sophisticated and successful strategy in some environments, may not guarantee high levels of ex-post allocative efficiency or distributional outcomes consistent with predictions of CE.

Thus, our results suggest that the famous Adam Smith quote about trade has a limitation. [29] A pure devotion to perceived self-interest in the marketplace results in a reduction of trading volume and resulting efficiency, because some, however slight, degree of benevolence towards one's fellow traders (in the form of a willingness to risk a mistake in adding liquidity to the market) is necessary to lubricate market trading.[30]

Although our results address the limited questions that could be tackled in controlled experiments, they are partially relevant as far as the general concerns regarding distributional outcomes are concerned. Our results need to be examined in alternative environments, e.g., in response to a variation of the one or more of the following: presence of trading strategies other than those explored here, the mixture over various trading strategies, the shape of the aggregate demand and supply schedules and other design parameters, to have a fuller understanding of the distributional properties of markets. Similarly, the use of simulated traders does not allow us to examine to what

[29] *" It is not from the benevolence of the butcher, the brewer, or the baker that we expect our dinner, but from their regard to their own interest. We address ourselves, not to their humanity but to their self-love, and never talk to them of our own necessities but of their advantages….."* Adam Smith (Wealth of Nations, Book 1, Chapter 2, 1776)

[30] After writing the above, we took notice of a paper by Vernon Smith renewing interest in Adam Smith's early writing, the *Theory of Moral Sentiments* and his caution that the famous Adam Smith quote in *Wealth of Nations* does not advocate purely selfish behavior in the market (Smith, V.L. (2013) Adam Smith: From Propriety and Sentiments to Property and Wealth, Forum *for Social Economics, DOI: 10.1080/07360932.2013.798241)*. If we understood the paper correctly – Adam Smith instead advocates a willingness to trade and communicate offers to a trading partner, something the random ZI traders in our study do more effectively than the Snipers who are waiting, looking for "great deals" that are better than should be expected rather than making their own offers. Part of this is what Vernon Smith calls "Smith's Discovery axiom" (p.11) within Adam Smith's works of "the propensity to truck, barter and exchange one thing for another" – or put more simply, people like to trade. Our robot traders (ZI or Sniper) lack the human quality to engage in complex behaviors such as 18th century era beneficence, judging propriety, or creating social codes of conduct. But to the extent that we can show an irrationally selfish and informationally stingy strategy can hurt a market overall, we think we are possibly in agreement with Vernon Smith's updated interpretation of Adam Smith's *Wealth of Nations.*

extent pro-social motivations crowd out self-interest driven motivations to influence behavior in DA markets that determines distributional outcomes. Further research can take up these challenges. These complementary endeavors are likely to help understand to what extent market processes contribute to observed inequality.

Finally, if the volume of transactions is used as a proxy for employment levels, the following observation by Stiglitz becomes pertinent. Stiglitz (2012) states that: "Unemployment-the inability of the market to generate jobs for many of its citizens- is the worst failure of the market, the greatest source of inefficiency, and a major cause of inequality." The results observed in our stylized markets where the volume of snipers was relatively high, any loss of allocative efficiency went hand in hand with more unequal distributional outcomes, are consistent with the observation made by Stiglitz that are derived from non-stylized settings. It seems that the distributional consequences of market processes need to be examined, more carefully, in environments which do not guarantee high levels of efficiency to discover any systematic links between reduced efficiency and distributional consequences.

References

Alton, M.R., and Plott, C.R. (2007). "Principles of Continuous Price Determination in an Experimental Environment with Flows of Random Arrivals and Departures." Working paper. Available at http://dx.doi.org/10.2139/ssrn.1083863

Becker, G. S. (1962). "Irrational behavior and economic theory." *Journal of Political Economy*, *70*(1), 1-13.

Bernanke, Ben S. (2007). "The Level and Distribution of Economic Well-Being," at https://www.federalreserve.gov/newsevents/speech/bernanke20070206a.htm

Brewer, P.J., Huang, M., Nelson, B., and Plott, C.R. (2002) "On the Behavioral Foundations of the Law of Supply and Demand: Human Convergence and Robot Randomness." *Experimental Economics* 5, 179-208.

Brewer, P.J., Cvitanic J., and Plott, C.R. (2013) "Market Microstructure Design and Flash Crashes: A Simulation Approach," *Journal of Applied Economics* 16(2), 2013, pp.223-250.

Cason, T.N. and Friedman, D. (1993). "An Empirical Analysis of Price Formation in Double Auction Markets." In D. Friedman and J. Rust (eds.), *The Double Auction Market: Institutions, Theories, and Evidence.* 253-283.

Chamberlain, E.H. (1948). "An Experimental Imperfect Market." *Journal of Political Economy* 56(2), 95-108.

Collins, Raymond F. (2017). *Wealth, Wages and the Wealthy: New Testament Insight for Preachers and Teachers.* Liturgical Press.

Delsol, J.P. (2017). "The Great Process of Equalization of Conditions," in Delsol, J.P., Lecaussin, N., and E. Martin *Anti-Piketty: Capital in the 21st Century*. CATO Institute, Washington, DC: 2017

Dragulescu, A. and Yakovenko, V.M. (2000). "Statistical mechanics of money" *The European Physical Journal B* 17, 723-729.

Eckel, C.C, Houser, D. Boettke, P.J. (2017). "Symposium in Honor of Vernon Smith: Introduction. A Celebration of Vernon Smith's 90th Birthday and Lifetime Contributions to Economics", *Southern Economic Journal* 83 (3), 639-643.

Easley, D., and Ledyard, J.O. (1993). "Theories of Price Formation and Exchange in Double Oral Auctions." In D. Friedman and J. Rust (eds.), *The Double Auction Market: Institutions, Theories, and Evidence.* 253-283.

Farmer, J. D., Patelli, P., & Zovko, I. I. (2005). "The Predictive Power of Zero Intelligence In Financial Markets." *Proceedings of the National Academy of Sciences of the United States of America*, *102*(6), 2254-2259.

Friedman, D. (1984). "On the efficiency of experimental double auction markets." *American Economic Review* 74(1), 60-72.

Friedman, D. (1993). "The double auction market institution: A survey." In D. Friedman and J. Rust (eds.), *The Double Auction Market: Institutions, Theories, and Evidence*. 3-25.

Giminez-Funes, E., Godo, L., Rodriquez-Aguilar, J.A., Garcia-Calves, P. (1998). “Designing bidding strategies for trading agents in electronic auctions.” In *Multi Agent Systems,* Proceedings of the IEEE, 136-143.

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, 119-137.

Hayek, F.A. (1945). “The Use of Knowledge in Society.” *American Economic Review* 35 (4), 519-530.

Kaplan, A. (1964). *The Conduct of Inquiry: Methodology for Behavioral Science*. Chandler Publishing Company.

Marshall, A. (1895) *Principles of Economics*, 3rd edition. London: Macmillian and Co. Available online at https://archive.org/details/principlesofecon01marsrich

Piketty, T. (1997). “The dynamics of the wealth distribution and the interest rate with credit rationing.” *Review of Economic Studies*, *64*(2), 173-189.

Piketty, T. (2014) *The Economics of Inequality*. Translated by A. Goldhammer. Harvard University Press: Cambridge.

Plott, C.R. and Smith, V.L. (1978) “An Experimental Examination of Two Exchange Institutions.” *Review of Economic Studies* 45 (1), 133-153.

Plott, C. R. (1982) “Industrial Organization Theory and Experimental Economics.” *Journal of Economic Literature* 20, 1495-1527.

Reiter, S. (1977). "Information and Performance in the (New) Welfare Economics." *American Economic Review 6*7(1), 226-234.

Rust, J., Miller, J.H., and Palmer, R. (1993). “Behavior of trading automata in a computerized double auction market.” In D. Friedman and J. Rust (eds.), *The Double Auction Market: Institutions, Theories, and Evidence.* 155-198.

Smith, A. (1950). *An Inquiry into the Nature and Causes of the Wealth of Nations, (1776)*. Methuen.

Smith, V.L. (1962). “An experimental study of competitive market behavior.” *Journal of Political Economy* 70, 111-137.

Smith, V.L. (1965). "Experimental Auction Markets and the Walrasian Hypothesis." *Journal of Political Economy* 73, 387-393.

Smith, V.L. (1982). “Microeconomic systems as an experimental science.” *American Economic Review* 72, 923-955.

Smith, V. L., and Williams, A. W. (1982). “The effects of rent asymmetries in experimental auction markets.” *Journal of Economic Behavior & Organization*, *3*(1), 99-116.

Smith, V.L., and Williams, A.W. (1990). "The Boundaries of Competitive Price Theory: Convergence Expectations and Transaction Costs." Green, L. and Kagel, J.H. (eds.) *Advances in Behavioral Economics, vol. 2*

Smith, V.L. (2013) Adam Smith: From Propriety and Sentiments to Property and Wealth, *Forum for Social Economics, DOI: 10.1080/07360932.2013.798241*

Stiglitz, J. E. (2012). *The price of inequality: How today's divided society endangers our future*. WW Norton & Company.

Walras, L. (1877). *Elements d'Economie Politique Pure*

Wilensky, U., and Rand, W. (2015). *An introduction to agent-based modeling: modeling natural, social, and engineered complex systems with NetLogo*. MIT Press.

Wilson (1984). “On equilibria of bid-ask market.” Working paper: Stanford University Institute for Mathematical Studies in the Social Sciences (No. TR-452).