Upending Stock Market Structure Using Secure Multi-Party Computation

Charanjit S. Jutla
IBM T. J. Watson Research Center
Yorktown Heights, NY 10598, USA
csjutla@us.ibm.com

Abstract

The stock markets have two primary functions, that of providing liquidity and price discovery. While the market micro-structure was mostly ignored or assumed to function ideally for the purpose of asset pricing, M. O'Hara (Journal of Finance, 2003) has established that both liquidity and price discovery affect asset pricing, and in particular asset returns. While the cost of liquidity provision is borne by investors, and is clearly detrimental to asset returns, periodic price discovery has both positive and negative consequences for asset pricing. In this work we propose using cryptography, and in particular multi-party secure computation, to setup a novel stock market structure that, to a large extent, removes the negative consequences of liquidity costs and periodic price discovery. Interestingly, the proposed market structure takes us back to the early days of stock markets, i.e. periodic call markets, but with the not so "trusted" auctioneer replaced by secure distributed computing where no individual party (or small coalition) gets to know the order book.

1 Introduction

While secure multi-party computation (MPC) [Yao86, GMW87, BOGW88, CCD88] has been adequately demonstrated to be useful in replacing auctioneers in clearing price auctions [BCD+09], its effectiveness in current stock market structures is questionable due to the high computational cost of MPC. Most of the current stock markets run a mechanism called (continuous) double auction (DA) where buyers and sellers continuously submit limit bids and asks, and any matching ask or bid is considered binding and a transaction takes place. With computerization, thousands, if not millions, of such trades can happen in a second in a single highly traded stock on NASDAQ. The limit bids and asks are in the clear (i.e. not sealed) at one or more central repositories, and traders can strategize based on this public information. The New York Stock exchange (NYSE) runs slightly differently in the sense that the limit bids and asks are with a "specialist" and the specialist keeps the bids and asks sealed (or secret from public), and is trusted to announce matching bids and asks as binding transactions. However, even NYSE of late has been slowly moving toward a NASDAQ style market.

Two natural questions arise. First, why have the stock markets moved toward continuous double auctions when clearing-price double-auctions (CPDA) seemed to be the norm in early stock

markets, and still are the norm in the U.S. treasury bond auctions and the opening round of NYSE. Second, should the limit bids and asks be in the clear in the double auctions?

The answer to both these questions may lie in the fact that the "specialist" is not generally trusted, especially for large orders. Another reason may be that stock market simulations show that continuous double auctions tend to converge to a Walrasian equilibrium [Wal77] (see e.g. [LS06]), and it even seems plausible to argue that the continuous double auction resembles Walras' tatonnement process (an almost end-less trial and error approach to "convergent" price-discovery; see e.g. [LS06]). However, it is also known that certain clearing-price double-auctions also converge rapidly to a Walrasian equilibrium, as long as the number of traders is large [SW89]. So, given that CPDA are equally efficient and the fact that Walras' tatonnement process really requires that the bids and asks in each round be secret, a fact which we will expound on later, the reasonable conclusion seems to be that the untrustworthiness of the specialist is the cause behind the trend toward non-specialist systems. To quote T. Cason and D. Friedman [CF97], "the benevolent auctioneer is at best an endangered species" (for empirical studies on negative specialist behavior, see [MC00, Ben06]).

Thus, if a distributed computerized process using cryptography, e.g. MPC, can emulate an "ideal" specialist or auctioneer then clearing-price double-auctions suddenly become more favorable. This is so, because unlike the open bid and ask DA which relies on the assumption that the DA mechanism somehow mimics Walrasian tatonnement, the CPDA can be shown to converge to Walrasian equilibrium under reasonable assumptions. We now argue that there is an even more important reason to favor sealed-bid clearing-price double-auctions.

Information about the economy is by nature decentralized. It is important to design mechanisms so that this information can be shared with others, who can then make better investment choices, for the "societal good". Hayek [Hay45] argues that it is not just scientific knowledge, but private knowledge about arbitrage opportunities, as well as inefficiencies in commerce that are important to be shared, and to be *not* assumed a priori given to all as is commonly assumed in standard economic theory. From his arguments one can conclude that the stock market's function is to garner this decentralized information into asset prices.

Taking a cue from Hayek's criticism of equilibrium theories of prices where all information is considered publicly known, Grossman and Stiglitz [GS80] further argue that the "efficient markets" hypothesis, i.e. that prices reflect all available information, implies that equilibrium does not exist or the market will vanish. In fact, using a Gaussian model of public and private information, they prove that if efficient markets hypothesis is true and prices reflect even the "informed" traders' private information, then the informed traders have no incentive to garner costly private information. Thus, assuming rationality, all traders will become "uninformed". But, this is also not an equilibrium, as clearly one trader can garner information, even at a fixed cost, and taking the "equilibrium" price as a given, increase his returns tremendously.

Building on [GS80], Easley and O'Hara [EO04] show that Walrasian equilibrium prices only partially reflect the informed traders knowledge about future prices, as long as the informed traders knowledge itself is Gaussian and they are risk-averse. This allows for the informed to continue to pay for costly research, as it allows them to be more competitive. Note, that in this Walrasian "pricing" economic model it is important that only prices, and distributions of various traders' information etc., be publicly known (and not the actual bid, ask price and quantity).

In this work, we show that a specific clearing price double auction mechanism that we describe incentivizes the informed trader even more. Further, when a large fraction of traders get

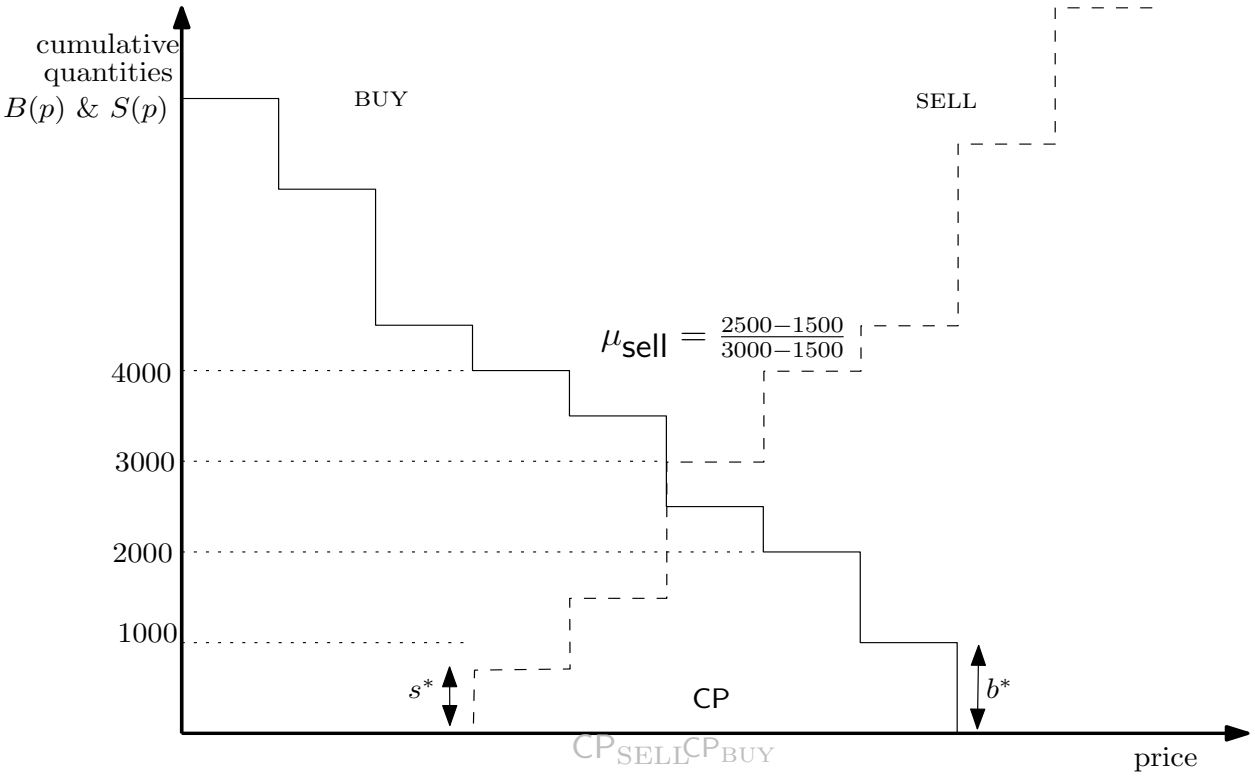

Figure 1: Clearing Price Double Auction

informed, either by eventual public dissemination of information (e.g. after an earnings report) or by expending resources on research, then the clearing price converges to Walrasian equilibrium prices. Moreover, even though the informed are further incentivized, as long as their knowledge is Gaussian, the uniformed still get returns by trading on the public knowledge of the informed's knowledge imprecision.

The specific clearing price double auction is similar to the usual CPDAs, except for two specific differences: (a) the traders are allowed to submit market price buy or sell orders, and the clearing price is such that all market price orders (buy and sell) are cleared, (b) if the price at which supply and demand cross, the supply and demand are not equal, then the clearing price is chosen so as to satisfy property (a), i.e. full clearance of market orders. This still sets the price to be either the greatest price where demand is larger than supply, or the least price where supply is larger than demand (see Fig. 1).

The rest of the paper is organized as follows. In Section 2 we describe the specific CPDA auction we propose as stock market mechanism. The security model is discussed in Section 2.6. The Walrasian equilibrium analysis of Easley and O'Hara in presence of private information ([EO04]) is reviewed in Section 3. Finally, the effects of micro-structure of our CPDA are illustrated in Sections 3.1 and 3.1.1.

2 Stock Market Mechanism

We first describe the mechanism assuming an ideal functionality $\mathcal{F}$ or an ideal auctioneer. An ideal functionality takes inputs from all the participants, computes the required function(s) and returns the specified outputs to the individual participants. No other information about the different participant's input is leaked to any of the participants. For an example of the following mechanism, see Fig. 1.

- All brokers submit their clients' bids and asks as follows:

  Each broker is pre-assigned a broker ID, say BROKER-ID.

  Each client can make two kinds of bids and two kinds of asks. One of the two kinds in either case is *market price*.

  If it is not a market price bid or ask, the bid or ask is a list of pairs $\langle z_i, p_i \rangle$, where z_i is the quantity bid to buy at price p_i (or quantity z_i to sell with ask-ing price p_i). Each price shall be in (positive integer) cents.

  If it is market-price the list of pairs is a singleton with the pair $\langle z, * \rangle$.

  Thus, each client's order will be specified as (order-type = market/limit, buy/sell, list-of-pairs).

  A client can submit multiple orders.

  The broker assigns a random ID to each order (large enough to be unique), say ID, and submits (BROKER-ID, ID, order-type = market/limit, buy/sell, list-of-pairs).

- The ideal functionality $\mathcal{F}$ computes the following:

  1. For each possible price p, $\mathcal{F}$ computes $b(p)$ as the sum of all z, where the sum is over all brokers and all of their orders marked BUY such that its corresponding list-of-pairs contains a pair $\langle z, p \rangle$.
  2. $\mathcal{F}$ also computes b^* as the sum of all z, where the sum is over all brokers and all of their orders marked BUY such that its corresponding list-of-pairs contains a pair $\langle z, * \rangle$.
  3. Similarly using the orders marked SELL, $\mathcal{F}$ computes $s(p)$ for each possible price p, as well as s^*.
  4. For each p, let $B(p) = b^* + \sum_{p' \geq p} b(p')$, and $S(p) = s^* + \sum_{p' \leq p} s(p')$.
  5. Next, $\mathcal{F}$ computes the least price p, denoted cp_{BUY}, such that $B(p) \leq S(p)$.
  6. $\mathcal{F}$ also computes the largest price p, denoted $\mathsf{cp}_{\text{SELL}}$, such that $B(p) \geq S(p)$.
  7. If $S(\mathsf{cp}_{\text{SELL}}) > B(\mathsf{cp}_{\text{BUY}})$, then clearing price $\mathsf{cp} = \mathsf{cp}_{\text{SELL}}$, otherwise $\mathsf{cp} = \mathsf{cp}_{\text{BUY}}$. Also, $\mathcal{F}$ outputs cp.
  8. All buy orders which are either market price or if their bidding price is greater than or equal to cp_{BUY} are given full allotment at price cp. $\mathcal{F}$ outputs all such bids, along with their BROKER-ID and ID.
  9. Similarly, all sell orders which are either market price or if their ask-ing price is less than or equal to $\mathsf{cp}_{\text{SELL}}$ are given full allotment at price cp. $\mathcal{F}$ outputs all such asks along with their BROKER-ID, ID.
  10. $\mathcal{F}$ also computes a fraction to a pre-determined precision.

– If $\mathsf{cp} = \mathsf{cp}_{\text{BUY}}$, $\mathcal{F}$ computes $\mu_{\text{SELL}} = (B(\mathsf{cp}_{\text{BUY}}) - S(\mathsf{cp}_{\text{SELL}}))/s(\mathsf{cp}_{\text{BUY}})$. Note, the denominator is the number of sell orders at exact price cp_{BUY}. We will assume that orders are split into groups of 1000 or lesser number of shares. $\mathcal{F}$ repeatedly picks one Such sell order (at limit price cp_{BUY}) at random, and outputs it with probability μ_{SELL} to be fully alloted at price cp. The process repeats as long as the sum total of such orders remains below $B(\mathsf{cp}_{\text{BUY}}) - S(\mathsf{cp}_{\text{SELL}})$.
– Else, $\mathcal{F}$ computes $\mu_{\text{BUY}} = (S(\mathsf{cp}_{\text{SELL}}) - B(\mathsf{cp}_{\text{BUY}})/b(\mathsf{cp}_{\text{SELL}})$. A similar randomized process selects and outputs the buy orders at price $\mathsf{cp}_{\text{SELL}}$ with probability μ_{BUY}.

Readers not familiar with secure multi-party computation are referred to [BCD+09] or [CHK+12][1]. These works describe different, but now standard, ways of implementing the ideal functionality using a distributed set of parties. In particular, the former is built on [BOGW88, CCD88] and the latter uses Rabin's Oblivious Transfer (OT) [JH85] and is built on [GMW87, NP05, IKNP03, LLXX05]. Both these protocols are in the semi-honest security model (see section 2.6 below), which means that the parties are assumed to not deviate from the protocol, but are allowed to collect any information they can from the execution of the protocol. In other words, assuming each party is participating in the secure multi-party computation protocol using a computer connected to the other computers via a network, then each party can observe and collect information from the memory of their computer during and after the execution of the protocol. Since the protocols are proved secure, this means that the parties cannot gather any additional information about the inputs of others by this semi-honest behavior, over and above what the ideal functionality will give them as output. While the second protocol above [CHK+12] also assures that no coalition of parties can gather any information beyond the outputs given to this coalition by the ideal functionality, the first protocol is only secure as long as there is no collusion. There is usually a performance trade-off when guaranteeing different security levels. Since, our functionality $\mathcal{F}$'s output is public (i.e. same for all participants), the first protocol [CHK+12] guarantees that as long as all the parties do not join hands, the inputs of the brokers remain secret (except for what can be inferred from the output of $\mathcal{F}$, or from a priori public information about the inputs).

Note that in the mechanism above, we assume that brokers run the protocol through $\mathcal{F}$. However, in the real world there can be a large number of brokers (from different countries, as well). Thus, to be practical, we will assume that some small subset of these brokers, say four, and a regulating authority such as the securities and exchange commission (SEC) actually run the protocol emulating and replacing $\mathcal{F}$. In standard cryptography parlance this would then be called a secure 5-PC (five party-computation). Recall, different 5-PC protocols give different security guarantees against collusion.

The brokers would still need to encrypt their bids and asks (for all its clients). For some versions of MPC, the brokers may alternatively be required to implement splitting their inputs using known secret sharing techniques and then sending the shares (using a cryptographically secure channel, e.g. SSL) to the five parties implementing MPC (the secret sharing can be a basic additive splitting or of a threshold kind [Sha79]). At the end of the protocol, the five parties will all know the output (as prescribed by $\mathcal{F}$), and they can publish it on a public bulletin-board. The actual process of doing the trade, i.e. which CUSIP numbered shares go to which broker can be done by some authority (it may charge a small fees for its services).

[1]For another MPC system which also compiles programs to circuits see [BDNP08].

2.1 How many CPDA sessions a day?

The number of CPDA auction sessions per day for each stock will depend on the time it takes to do the 5-PC. The opening session of each day usually commands the most attention, and it may have the most volume of inputs, and hence is likely to take thirty minutes to compute (by current computational power of AMD/Intel servers). The other sessions may complete in 15 mins. One must also provide a gap of about five minutes between sessions for all clients to digest the results of the previous auction. Note that after that five minute time gap, no more inputs can be provided to that session.

Remark. We remark that repeated CPDA sessions are still prevalent with a *specialist* acting as the "trusted" auctioneer. For instance, the Taiwan commodities market only recently shifted from such a repetitive call market to a continuous (double) auction market [KL11]. This also served as a useful case study of comparing the two mechanisms and [KL11] conclude that "while the quoted spreads, effective spreads, and price volatility are all smaller in the continuous auction market, the call auction market exhibits greater market depth and smaller pricing errors; the latter is also found to be more effective in resolving the problem of information asymmetry".

2.2 Low Volume Market Price

Usually during the main trading sessions, the fraction of trades at market price to those with limit is low enough that the above methodology is sound. However, a caveat can be introduced that if $\min\{s^*, b^*\}$ is more than $\min\{S(\mathsf{cp}), B(\mathsf{cp})\}$, the the clearing price is reset to the clearing price of the previous session; in other words, in such a case the market price is just the clearing price of the previous session. If $B(\mathsf{cp}) < S(\mathsf{cp})$, then all market sell orders and sell order with limit price less than or equal to cp are executed, and only an appropriate fraction of the buy orders are executed. Similarly, for the opposite case.

Another possibility is to let the market price orders be in the clear, so that competitive forces can provide a better price if the non market-price volume is low.

2.3 Upstairs Market

The New York Stock Exchange consists primarily of the "downstairs" market where a specialist handles and clears all bids and asks. However, many liquidity providers (other than the specialist) are reluctant to leave large limit orders in the specialist's order book, as they would then run the risk of trading with someone with new information. This means that an uninformed seller who is just seeking liquidity of a large chunk of shares (say, to diversify into some other investment) may not get it in the "downstairs market". However, such traders can be advised to go "upstairs" where the specialist firm will help the trader find other traders to take the opposite side of the trade if the latter can be convinced that the former does not have any new information. Essentially, liquidity providers are reluctant to trade with some known traders who are well-known to have additional information [MC00].

Clearly in our mechanism, such liquidity providers would again be reluctant to leave large limit orders. However, the upstairs market can still exist with their usual commission for their services. The two parties, after having being convinced of the former's identity, can place their large matching limit orders in the CPDA. With high probability, the orders will go through at that price, or at the very least the former trader's order will be executed at his/her limit price or better.

2.4 Improved Liquidity

Regardless of whether the "upstairs market" continues to be used in our system or not, the fact that the limit bids and asks are now completely secret (other than what is disclosed from the output of $\mathcal{F}$, and assuming no collusion among the five parties beyond the security provision of the particular 5-PC), the main concern of the traders in submitting large bid and ask goes away. For example, a typical example often cited in the current stock market structure is that not all investors maybe present at the same time, even though they are willing to take opposite sides of the trade. An investor A may want to sell 1000,000 shares of Apple stock today at current prices, whereas another investor B may want to buy 1000,000 shares at current prices, but only tomorrow due to some independent reasons (assuming no new drastic information is generated). In the current system, if a large sell order by A is put into the market, the prices immediately drop (this behavior of the market holds even in the NYSE specialist system [MC00]). In essence, investor A ends up paying the cost of liquidity to other intermediaries who serve as liquidity providers. However, in our system, as long as investor A splits his million shares into 1000 blocks of 1000 limit orders each (and is patient for a day or two, or sometimes even a few CPDA sessions), partial information leaked by outputs of $\mathcal{F}$ in different CPDA sessions, will still hide the remaining open orders of A. We model and analyze this behavior of our mechanism more rigorously in Section 3.1.1.

It has often been cited [BHR13] that high frequency trading (HFT) has improved liquidity in the DA markets. However, it has also come under severe anecdotal criticism [Lew14] that the HFT traders may completely remove liquidity when the market is under stress. Moreover, the improved liquidity possibly comes at a higher liquidity cost, as all the bids and asks are open in the DA markets (note that once the bids and asks are in the open, the analysis of [EO04] which shows that informed traders have an incentive does not hold any more). Our mechanism, which is closer in spirit to the specialist system (without the concern of specialist leaking the order book), is expected to remedy these HFT front-running and high liquidity-cost concerns.

2.5 Improved Incentives for Information Gathering

We also show in Section 3.1.2 that our mechanism incentivizes the informed trader even more than as predicted by the Walrasian equilibrium analysis of [EO04]. However, the uninformed traders continue to get returns as the information gathering is not perfect. Further, eventually all information becomes more or less public, e.g. after an earnings report, and the playing field is leveled for rational traders. Thus, the uninformed (but rational) traders continue to be in the market and be the usual source of liquidity.

2.6 Security Model: Semi-Honest vs Malicious Adversary

While secure multi-party computation can also be achieved under malicious adversarial behavior, e.g. by using verifiable secret sharing [CGMA85, JMN10], in the stock market setting this may not be required if all the participants are required to save all their computations and release them to the public after a month or so. The SEC (or any party for that matter) can audit the saved computations and check if all computations were performed according to protocol. Note, in our economic modeling we do not require that the bidding strategies of the traders are secret. In fact, thees strategies, and even all distributions are assumed to be public knowledge. Thus, if the transcripts are released after a reasonable amount of time, there is no harm to the incentives that the informed traders are receiving in our system.

3 Modeling Information Risk and Research Advantage

Since double auctions are notoriously difficult to analyze, we will make some simplifying assumptions. We will assume that the seller(s) have a supply of stocks to sell, which follows a normal distribution. In the simplest case, sellers just ask market price for their supply of shares. A slightly more complicated case will be considered below.

Each buyer has a choice of investments which can either be cash, or one of the n risky stocks. Let's say at the beginning of a day when the buyer needs to make an investment choice, it is public knowledge that the value of the k-th stock in the future (or next day), denoted v_k, will follow a normal distribution with mean $\bar{v}_k$ and precision ρ_k. This, of course, is a base case distribution. Some traders may have proprietary research (which could be based on public information such as filings by the company or it could be based on private research and modeling) that could change their future valuation of each stock. For simplicity, we will assume that such *informed traders* receive the same signal s_k, and the probability distribution of s_k is itself a normal distribution with mean v_k and precision γ_k conditioned on the future valuation of the stock being v_k. Thus, s_k more or less predicts v_k with precision γ_k. In fact as observed in [EO04], conditioned on the signal s_k, the distribution of v_k can be calculated by Bayes rule to be a normal distribution with mean $\bar{v}_k\rho_k + s_k\gamma_k$ and precision $\rho_k + \gamma_k$.

We will also assume that each buyer (trader) T_j has an exponential utility function with risk-aversion coefficient δ^i. In other words, the buyer's utility function u_j is the following function of wealth w_j: $u_j(w_j) = (1-e^{-\delta_j w_j})/\delta_j$. It is safe to assume that big institutions are more risk-neutral, meaning their δ is close to zero, whereas individual investors will tend to have a large δ.

For any buyer T_j with risk-aversion coefficient δ_j, and initial wealth w_j, suppose T_j buys $x_{j,k}$ quantity of the k-th stock at price p_k. Then, the cash d_j it has left is $w_j - \sum x_{j,k}p_k$. Thus, their future wealth Ψ_j is $d_j + \sum x_{j,k}v_k$, which is same as

$$w_j + \sum_k (v_k - p_k) \cdot x_{j,k}$$

Now, suppose that the buyers know the price they have to pay for the k-th stock, e.g. price in Walrasian equilibrium [Wal77, LS06]. Then, their future utility can be maximized w.r.t. each of the quantities $x_{j,k}$. For the buyers who receive a signal for the k-th stock, we can also condition on the signal being s_k. In either case, i.e. whether they receive a signal or not, the distribution of v_k remains normal with a mean and precision known to the buyer. As mentioned earlier, if they do receive a signal s_k, then the conditional distribution of v_k is still normal with mean $\bar{v}_k\rho_k + s_k\gamma_k$ and precision $\rho_k + \gamma_k$.

It can be shown that if the utility function of T_j is of the above exponential form with risk-aversion coefficient δ_j, then it suffices to maximize $\mathrm{E}[\Psi_j] - (\delta_j/2) \cdot \mathrm{Var}[\Psi_j]$. It is not difficult to see that maximum is achieved at

$$x_{j,k} = \frac{\bar{v}_{j,k} - p_k}{\delta_j(\rho_{j,k})^{-1}}$$

where $\bar{v}_{j,k}$ and $\rho_{j,k}$ are the (conditional) expectation and variance resp. of v_k (conditioned on the signal s_k for the buyers who receive signal s_k).

With this analysis, [EO04] conclude that if *all* traders receive the signal, and assuming all of them have the same δ, and further if there is a per capita supply y_k of the k-th stock, then since

in Walrasian equilibrium supply equals demand, we get that

$$p_k = \frac{\bar{v}_k \cdot \rho_k + s_k \cdot \gamma_k - \delta \cdot y_k}{\rho_k + \gamma_k}.$$

However, the situation gets more interesting if only a faction μ_k of traders get the signal for the k-th stock. In the Walrasian equilibrium model, the uninformed buyers (i.e. those who do not receive the signal) can estimate the signal from the converging price p_k (e.g. by a seemingly endless trial and error process known as "tatonnement"), as long as the distribution of the signal is public knowledge. In particular, [EO04] show that in case $\mu_k > 0$ then conditioned on the price p_k, the uninformed buyers can compute a random variable θ which is distributed normally with mean v_k (not to be confused with $\bar{v}_k$) and precision ρ_θ given by

$$\left[\gamma_k^{-1} + \left(\frac{\delta}{\mu_k \gamma_k}\right)^2 \eta_k^{-1}\right]^{-1}$$

Note that if δ is zero, as is the case for risk-neutral buyers, then this is the same distribution as that of the signal, and hence the price completely reveals whatever advantage the informed buyers get from the signal. This, however brings back the critique of Grossman and Stiglitz that in such an equilibrium the informed have no incentive to do research; or in other words, if they did support costly research, then no equilibrium is possible. However, if δ is non-zero, then the price is only partially revealing to the uninformed in equilibrium, and there remains an advantage to the informed. However, in the next section we show that the sealed bid clearing price auction as described above allows for the informed to retain an advantage even when they are risk-neutral.

Similarly, if the stock supply is accurately known publicly, i.e. η_k is infinity (e.g. in an initial public offering), then again the advantage of the informed goes away in the above Walrasian equilibrium. Fortunately, the mechanism we consider is not completely efficient (w.r.t. Walrasian equilibrium) .

3.1 Effects of Micro-Structure of Price Discovery on Asset Returns

While the analysis above was done assuming Walrasian demand-supply equilibrium without regard to the mechanism which may be required to achieve such an equilibrium, we now focus on practical mechanisms and analyze the effects on returns for both informed and uninformed traders. Again, since double auctions are difficult to analyze, we will focus on some simple cases, especially from the point of sellers. The results are more illustrative than being comprehensive.

The main goal is to show that sealed-bid and sealed-ask clearing price double auctions of the particular form as described above have the property that (a) the average return of an informed investor is higher than in the basic Walrasian equilibrium analysis, and (b) if all investors are equally informed then the clearing price rapidly converges to the Walrasian equilibrium price.

Actually, the property (b) has been well-studied in the individual private belief model, and it is shown in [SW89] that if the number of traders is large then they converge to bidding truthfully, i.e. in accordance with their private valuation. Also see [KN04], where it is shown that discrete bidding strategies (as opposed to continuous bidding functions) make the clearing price auction efficient.

So focusing on (a), we first remark that rewarding an investor who has put costly effort into researching a stock is not just important from an equilibrium perspective, but one can reasonably argue that such an equilibrium tends to incentivize research and hence more efficient utilization of

capital. Further, the uninformed also continue to have positive returns, since the signal that the informed receive is not absolutely precise. Thus, the uninformed remain in the game, and hence provide liquidity. Finally, when the news or information is public, e.g. after an earnings report, the playing field is leveled and the uninformed get even higher returns.

3.1.1 Price Discovery under Common Public Information

We first focus on the case where no buyer receives a signal (and it is common knowledge), and the distribution of the future price v is public knowledge to be a normal distribution with mean $\bar{v}$ and precision ρ. All buyers are assumed to have the same risk-aversion coefficient δ. We will also assume that there is a single seller, or the sellers know the total supply of stock for sale[2]. Assume that Δ is the minimum precision of bid and ask limit prices. Then, we claim that the following strategies for the buyers and sellers converges to a Nash equilibrium. For each price $p = \pi \cdot \Delta < \bar{v}$, each buyer puts in an order to buy $\Delta\rho/\delta$ shares of stock at limit price p. If there are N buyers, this means that $B(\pi \cdot \Delta) = N \cdot (\bar{v} - \pi \cdot \Delta) \cdot \rho/\delta$. The seller(s) puts in a sell order of $N \cdot x$ shares of its supply at limit price $\bar{v} - \rho^{-1}\delta \cdot x$.

With these strategies, if the total supply of shares is $N \cdot x$, the clearing price is $p = \bar{v} - \rho^{-1}\delta \cdot x$ (as Δ tends to zero). At that clearing price $B(p) = N \cdot (\bar{v} - p) \cdot \rho/\delta$, and each buyer gets exactly the number of shares as found in a Walrasian equilibrium. Given that other buyers and the seller play the above strategies, no single buyer can improve his utility by playing differently (at a higher price, his demand is lower, and a lower price bid by him will not clear). At is a simple exercise to see that the sellers can even put in a market order and achieve Nash equilibrium as long as the number of buyers tends to infinity.

The situation is more tricky, if there are multiple buyers and then total supply of stock for sale is not known to any individual seller. As mentioned earlier, if there are a large number of buyers, then the sellers can even try market order. But it is not immediately clear if one seller can get an advantage by asking a limit price. Also, the situation gets really tricky if the the sellers are not risk neutral; the greater the supply, and more the price drops, they may be incentivized to hold on to some supply. Another interesting scenario is where the buyers are not all similarly risk-averse. The analysis of such situations merits further study.

3.1.2 Price Discovery under Mixed Private and Public Information

The situation gets really interesting when only a fraction of buyers receive a pointed signal about the future stock price(s). Recall that in such a scenario, in the Walrasian equilibrium setting above (section 3), the price is only partially revealing to the uninformed buyers. We will now illustrate that in the CPDA considered in this work, that the informed are even more incentivized to gather information than suggested by the Walrasian equilibrium (and are incentivized even if they are risk neutral). Moreover, the uninformed also continue to get a positive return on their investment.

Remark. The argument that the uninformed will continue to invest in stocks in this market mechanism holds water only if it can be demonstrated that they are better off in this system vs. the current continuous DA system. To this end, we need to make a further distinction between traders. We will classify traders as *informed*, *uninformed-but-curious* and *uninformed-but-rational*. The

[2]The sellers are risk-neutral ($\delta = 0$), and have found a better investment than this stock; thus, selling their supply of stock is in their best interest.

uninformed-but-curious do not expend resources on researching the stock, but will try to gain an advantage from the inefficiency of the particular market mechanism. This should be contrasted with Hayek's arbitrageurs, who have specialized intelligence on arbitrage in the economy as a whole, and hence are to be categorized as informed. However, the uninformed-but-curious can provide a useful service by providing liquidity to traders in need of urgent liquidity. Such service can be provided for a commission and/or by a negotiated price as in the upstairs market (see 2.3). The uninformed-but-rational do not do any research on the stock(s), nor are they mathematically inclined to figure out conditional distributions etc. In a nutshell, they are rational in the sense that they known that $\bar{v}$ is greater than the current prevailing price (say, off the last few CPDA sessions) and they will place a market order knowing well that the quantity of market orders will be much smaller than limit orders placed by informed traders. Moreover, it is likely that if the informed are highly incentivized in our system, the uninformed will try to join the ranks of the informed, e.g. by paying a fees for information or by investing in mutual funds. Essentially, the only traders who lose in our system (compared to open bid-ask continuous DA) are the uninformed-but-curious.

For simplicity, we will assume that there are exactly two informed buyers, but with a much smaller risk-aversion coefficient δ'. Assume there are $N/2$ other uninformed buyers with risk-aversion coefficient $\delta >> \delta'$. In fact, we will assume that $N \cdot \delta' = 4\delta$. It will also be useful to assume that $\gamma >> \rho$, i.e. the signals are much more accurate than general information about v. The sellers are assumed to have a supply of shares to sell with a normal distribution with mean $\bar{y}$ and variance η. As opposed to the case where no buyer receives a signal, here the cumulative demand function is not publicly known. For this reason, the sellers, who are assumed to rid themselves of their supply, can be seen to follow the market price strategy at best. Further, since there are at least two informed traders, Nash equilibrium requires them to bid as follows: if the signal they receive is s (from a normal distribution with mean v and precision γ), then for each price $p = \pi \cdot \Delta$, such that $p < (\bar{v}\rho + s\gamma)/(\rho+\gamma)$, they put in a limit order for amount $\Delta \cdot (\rho+\gamma)/\delta'$. The uninformed continue to follow the usual strategy: for each price $p < \bar{v}$ then put in a limit order for amount $\Delta \cdot \rho/\delta$.

Since the uninformed can not base their bids on the partially revealing Walrasian equilibrium price of Section 3, the informed now get a better return. As for the uninformed, if the signal s is much larger than v, and the volume of stock for sale is not high enough, they may be priced out, and get zero allocation. Also, when the supply of stock is much below average and s is much below $\bar{v}$, the uninformed are likely to get full allocation when it is sub-optimal (e.g. the informed do not even bid). However, the probability of this combined event is very low. Otherwise, they continue to get a return as before (i.e. as in equilibrium pricing). The most to lose are the sellers (compared to when there was no signal for any trader). Indeed, with the assumption we have on the sellers, it is not unreasonable to assume that these uninformed sellers may have accumulated shares from previous rounds in tatonnement (or even worse in open bid continuous DA).

Remark. Note that in our modeling, at the end of a round of CPDA, both the price and volume information is revealed. In the simplistic modeling we considered above, this completely reveals the signal s (with high probability). However, this is only possible if the number of informed traders as well as the number of uninformed traders (as well as their risk profiles) was accurately known. Moreover, we have approximated the model further by assuming a market order strategy by the sellers. Further, in a more realistic setting, one must consider different kinds of signals (i.e. with

different precisions). By central limit theorem, one may assume that they still follows a public normal distribution with mean $\bar{v}$ and precision γ, but now the informed traders may not just get a signal s, but may also privately get more precise variance of this signal. Even if we ignore these more precise modeling issues and stick to our simple model, for this particular round of CPDA the informed are incentivized much more than in Walrasian equilibrium. Thus, at the very least, it supports the idea of using CPDA to open daily trading, as is done in NYSE, but with secure MPC replacing the specialist. Finally, we remark that the ideal functionality $\mathcal{F}$ in Section 3 can also try to hide the volume information, or only reveal the order of volume (say in powers of tens), by also matching buyers and sellers (by sending a message to matching clients). However, in view of the above arguments, we consider this unnecessary.

References

[BCD+09] Peter Bogetoft, Dan Lund Christensen, Ivan Damgård, Martin Geisler, Thomas P. Jakobsen, Mikkel Krøigaard, Janus Dam Nielsen, Jesper Buus Nielsen, Kurt Nielsen, Jakob Pagter, Michael I. Schwartzbach, and Tomas Toft. Secure multiparty computation goes live. In *Financial Cryptography and Data Security, 13th International Conference, FC 2009, Accra Beach, Barbados, February 23-26, 2009. Revised Selected Papers*, pages 325–343, 2009. 1, 2

[BDNP08] Assaf Ben-David, Noam Nisan, and Benny Pinkas. FairplayMP: a system for secure multi-party computation. In Peng Ning, Paul F. Syverson, and Somesh Jha, editors, *ACM CCS 08*, pages 257–266. ACM Press, October 2008. 1

[Ben06] Sigridur Benediktsdottir. An empirical analysis of specialist trading behavior at the new york stock exchange. *FRB International Finance Discussion Paper*, (876), 2006. 1

[BHR13] Jonathan Brogaard, Terrence Hendershott, and Ryan Riordan. High frequency trading and price discovery. Working Paper Series 1602, European Central Bank, 2013. 2.4

[BOGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *20th ACM STOC*, pages 1–10. ACM Press, May 1988. 1, 2

[CCD88] David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols (extended abstract). In *20th ACM STOC*, pages 11–19. ACM Press, May 1988. 1, 2

[CF97] Timothy N Cason and Daniel Friedman. Price formation in single call markets. *Econometrica: Journal of the Econometric Society*, pages 311–345, 1997. 1

[CGMA85] Benny Chor, Shafi Goldwasser, Silvio Micali, and Baruch Awerbuch. Verifiable secret sharing and achieving simultaneity in the presence of faults (extended abstract). In *26th FOCS*, pages 383–395. IEEE Computer Society Press, October 1985. 2.6

[CHK+12] Seung Geol Choi, Kyung-Wook Hwang, Jonathan Katz, Tal Malkin, and Dan Rubenstein. Secure multi-party computation of Boolean circuits with applications to privacy in on-line marketplaces. In Orr Dunkelman, editor, *CT-RSA 2012*, volume 7178 of *LNCS*, pages 416–432. Springer, February / March 2012. 2

[EO04] David Easley and Maureen O'hara. Information and the cost of capital. *The Journal of Finance*, 59(4):1553–1583, 2004. 1, 2.4, 2.5, 3

[GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th ACM STOC*, pages 218–229. ACM Press, May 1987. 1, 2

[GS80] Sanford J Grossman and Joseph E Stiglitz. On the impossibility of informationally efficient markets. *The American Economic Review*, pages 393–408, 1980. 1

[Hay45] F. A. Hayek. The use of knowledge in society. *American Economic Review*, 35(4), 1945. 1

[IKNP03] Yuval Ishai, Joe Kilian, Kobbi Nissim, and Erez Petrank. Extending oblivious transfers efficiently. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 145–161. Springer, August 2003. 2

[JH85] M.O. Rabin J. Halpern. A logic to reason about likehood. In *Proc. 15th ACM STOC*, pages 363–365, 1985. 2

[JMN10] Thomas P. Jakobsen, Marc X. Makkes, and Janus Dam Nielsen. Efficient implementation of the Orlandi protocol. In Jianying Zhou and Moti Yung, editors, *ACNS 10*, volume 6123 of *LNCS*, pages 255–272. Springer, June 2010. 2.6

[KL11] Weiyu Kuo and Yu-Ching Li. Trading mechanisms and market quality: Call markets versus continuous auction markets. *International Review of Finance*, 11(4):417–444, 2011. 2.1

[KN04] Ilan Kremer and Kjell G Nyborg. Underpricing and market power in uniform price auctions. *Review of Financial Studies*, 17(3):849–877, 2004. 3.1

[Lew14] M. Lewis. *Flash Boys: A Wall Street Revolt*. Penguin, UK, 2014. 2.4

[LLXX05] Bao Li, Hongda Li, Guangwu Xu, and Haixia Xu. Efficient reduction of 1 out of n oblivious transfers in random oracle model. Cryptology ePrint Archive, Report 2005/279, 2005. `http://eprint.iacr.org/2005/279`. 2

[LS06] Jonathan Levin and Ilya Segal. General equilibrium, 2006. `http://web.stanford.edu/~jdlevin/teaching.html`. 1, 3

[MC00] A. Madhavan and M. Cheng. Price discovery in auction markets: A look inside the black box. *The Review of Financial Studies*, 13(3):627–658, 2000. 1, 2.3, 2.4

[NP05] Moni Naor and Benny Pinkas. Computationally secure oblivious transfer. *Journal of Cryptology*, 18(1):1–35, January 2005. 2

[Sha79] A. Shamir. How to share a secret. *Comm. ACM*, 22(11):612–613, 1979. 2

[SW89] Mark A. Satterthwaite and Steven R. Williams. The rate of convergence to efficiency in the buyer's bid double auction as the market becomes large. *The Review of Economic Studies*, 56(4):pp. 477–498, Oct., 1989. 1, 3.1

[Wal77] L. Walras. *Elements of Pure Economics*. 1877. 1, 3

[Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167. IEEE Computer Society Press, October 1986. 1