Everything you wanted to know about double auctions, but were afraid to (bid or) ask

Simon Parsons
Department of Computer & Information Science
Brooklyn College
City University of New York
2900 Bedford Avenue, Brooklyn,
NY 11210, USA

Marek Marcinkiewicz and Jinzhong Niu
Department of Computer Science,
Graduate Center,
City University of New York,
365 Fifth Avenue, New York,
NY 10016, USA

Steve Phelps
Department of Computer Science
University of Liverpool
Chadwick Building
Peach Street, Liverpool
L69 7ZF, UK

August 18, 2006

Abstract

With the rise of ecommerce and the increase in volume of business procurement carried out electronically, double auctions have received much attention. This attention has been largely in the area of economics, but, increasingly, is in the area of computer science. However, there is no such thing as "the double auction". Instead there is a profusion of terminology and a wide variety of different market institutions that are all, confusingly, described as "double auctions". This paper aims to identify some order in this profusion. First, the paper aims to document the terminology and to identify the variety of market institutions that are grouped under the "double auction" heading. Second, the paper describes some of the theoretical and experimental work that has been carried out on double auctions. Finally, the paper sketches some of our recent work which

aims to synthesize a common set of experiments from those detailed in the literature.

1 Introduction

In recent years, double auctions have become a topic of much interest in both the economics and computer science literature. While work from the economics side can be seen as part of the steady progress of auction theory from simpler kinds of market, such as Dutch auctions, to more complex markets, that on the computer science side seems to be driven by the growth of ecommerce. With many new electronic markets starting up, computer scientists are increasingly involved in building market systems. Furthermore, market-based techniques [6, 71] are finding application as general approaches to resource allocation, often through the creation of artificial (in the sense that they do not involve the exchange of money) markets.

From some of the descriptions of the double auction institution that one often finds in the literature describing this recent work, it is easy to get the impression that "the double auction" is a single concept like the English or Dutch auction[1]. However, on looking more closely into the matter (for example when trying to implement such a market), it quickly becomes clear that there is a wide range of different variants of the double auction. The first aim of this paper is to lay out some of these variations, and, in particular, the differences between them. The second aim of this paper is to describe some of the theoretical and experimental work that has been carried out on double auctions, some of which has been unjustly ignored by a computer science audience, and to identify what this work does and does not tell us (with an eye to spotting places where additional experimental work is necessary).

It should be noted that this paper concentrates on a specific segment of the literature on double auctions. It focusses on work that is concerned with abstract markets — typically simple commodity markets — rather than the detail of specific markets (though it does describe the operation of the New York Stock Exchange at some length), on work that examines the equilibrium properties of such markets, and, especially, on work that deals with how traders make descsisions on how to make offers. It is not, for example, interested in the kind of detailed modelling of financial markets explored in [50] or [**?**]. This interest reflects our wider interest — which is in the design of new traders and the interaction between such designs and the designs of markets rather than in understanding existing markets.

1.1 Terminology

We start by considering the description of the double auction market given by Friedman [14], which kicks off the influential volume [15] describing work from

[1] Though there are variants of both the English and Dutch auctions, the range of this variation is small.

the Sante Fe Insitute conferences on double auctions that were held in 1990 and 1991.

We deal with interactions between agents, *traders*, and commodities that they own. We will usually deal with two commodities, the *good* and *money*. Sometimes this will be extended with additional kinds of good. All goods are taken to be indivisible, while money is divisible. The interactions between the agents lead to *exchange*, a process whereby the traders freely alter the allocation of commodities, without changing the total quantity of the commodities. Thus traders may choose to swap a good for some quantity of money, and the amount they choose is related to the value the trader places on the good. This value is known as the trader's *private value* or *limit price*. The latter terminology comes from the fact that if a buyer pays more than its private value, or a seller accepts less, then it considers it is trading at a loss.

A *market institution* defines how this exchange takes place. It does this by laying down rules about what the traders can do (which comes down to specifying what messages they can exchange) and rules for how the final allocation of commodities given the actions of the traders. When the allocation of commodities changes, a process known as the market *clearing*, the difference between the final and initial allocations is the *net trade*, and this has a component for each trader in the market which shows the change in the amount of money and good that they hold. The minimal net trade is when two traders have non-zero components, and any trader that has a non-zero component has an associated *transaction price* (which we will also sometimes call the *trade price*)—the absolute value of the money component divided by the good component[2]. Traders with positive good components are *buyers*, those with negative good components are *sellers*. Typically we will deal with so-called one-way traders, meaning that traders are either buyers or sellers but not both.

So far we have not said much about what form the market institution takes. Friedman [14] defines an *auction* as a market institution in which messages from traders include some price information—this information may be an offer to buy at a given price, in the case of a *bid*, or an offer to sell at a given price, in the case of an *ask*—and which gives priority to higher bids and lower asks. We can allow only buyers or only sellers to make offers, in which case the market is *one-sided*, or we can allow both, in which case it is *two-sided*. A double auction is a two-sided auction, and from here on we will only deal with double auctions (though many of the distinctions we identify could apply to other kinds of market as well)[3].

1.2 Structure

Now, given this set of definitions, we are ready to proceed with the cateloging of the double auction literature. The structure of this paper is as follows. Section 2 identifies the varieties of market that one can find in the literature, and Section 3

[2]Neither component is allowed to be zero if the other is zero.

[3]Friedman [14] distinguishes *quasi-auctions* as well as other kinds of market—a quasi-auction has most of the features of an auction but not all.

describes in detail one particular double auction market, the New York Stock Exchange. Section 4 discusses the theoretical work on some of the varieties of double auction, and Section 5 then briefly surveys experimental work on the double auction, leaning heavily towards work that uses artificial agents rather than human traders. Finally, Section **??** describes some of the work we have been undertaking which aims to replicate and relate work described in previous sections. Finally Section 6 concludes.

2 Types of double auction

The aim of this section is to suggest a framework into which all existing work on double auctions can be fitted, and to then identify into which part of the framework that existing work fits. This suggested framework is not really new, but a synthesis of suggested terminology and partial clasifications that are scattered across the literature.

2.1 Attributes of double auctions

A double auction can be *one-shot* or it can be *repeated*. If it is repeated (and as we will see below, most experimental work has looked at repeated auctions), then the auction has several, as opposed to one, *trading periods*. The reason that it is worth distinguishing repeated auctions as opposed to one-shot auctions is that in repeated at the start of each trading period agents receive a new allocation of commodities thus modelling [61]:

> conditions of normal supply and demand

but have the scope to remember the results of trading in the previous period. Thus, after several periods, we can expect the market to reach some kind of steady state. In this way trading periods often correspond to calendar days, with traders receiving new allocations at the start of each day, and so the term "day" is often used as a synonym for "period". Of course, running a number of periods with the same supply and demand curve is a nicety of experimental work — in real auctions supply and demand curves are not static, and some work has considered such cases. For example [24] and [41] look at changes in supply and demand in the context of the New York Stock Exchange, and [19] consider the performance of artificial trading agents under such conditions.

The notion of periods can easily be confused with what we will call *rounds* (known under several names in the literature). Often in experimental work, bidding between traders is synchronised in rounds — in each round every trader has the chance to make an offer. This is typically an artifact of an experimental auction between automated traders, but one can imagine real double auctions set up in this way, with the auctioneer controlling the submission of offers (but not their content). There may be many rounds in a period, and the usual situation in the literature seems to be to have sufficient rounds that every trader has the chance to trade if they desire to.

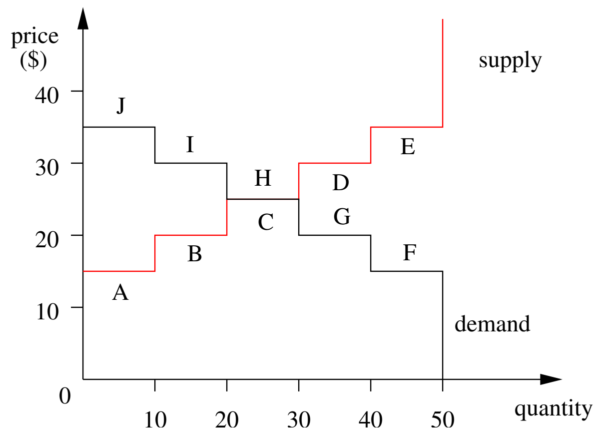

Figure 1: Illustrative supply and demand curves for a double auction.

In a *periodic*[4] or *discrete time* double auction, clearing happens some fixed period of time after the auction starts. This clearing might happen several times in each trading period, or may happen just once at the end of the trading period. The classic example of a discrete time double auction is the *call market* or *clearing house*[5] used to determine the opening prices on the New York Stock Exchange. Traders submit bids and asks until the end of the trading period and these are used to determine the supply and demand curves for commodities. Example supply and demand curves, taken from [6], are shown in Figure 1. These are based on having traders A, B, C, D and E all looking to supply 10 goods with limit prices of \$15, \$20, \$25, \$30, and \$35 per good respectively, and F, G, H , I, and J looking to buy with limit prices of \$15, \$20, \$25, \$30, and \$35 respectively. If the trade price is set below \$15, no seller will be prepared to trade, while if the price is between \$15 and \$20 then only A will trade, and so on. Similarly, at a price above \$35 no buyer will trade, while at prices between \$30 and \$35, only J will trade, and so on.

Once these curves are established, they are used to set a price for trading, the price at which supply equals demand in this case[6]. The price at which supply equals demand is known as the *equilibrium price*, and in the case of the market in Figure 1 is a price of \$25. In this kind of auction, it is an insitution external to the traders—the clearing house that gives the market one of its names—that is the entity that decides the trade price. In other

[4] "Periodic" in this sense does not, confusingly, relate to whether an auction operates over one or many periods in the sense described above. As a result, we will prefer the term "discrete time" to "periodic" in this paper.

[5] Though "call market" and "clearing house" seem to be the most common names for a discrete time market, [38] also class such entities as "batch markets". Note that Friedman [14] uses "call market" to indicate a market that clears just once, while a "clearing house" is a market that clears many times.

[6] Though as [40] explains, as as we will discuss below, the New York Stock Exchange doesn't work quite like this.

kinds of auction the traders themselves set the price, and we will distinguish the two cases as *institutional price-setting* and *non-institutional price-setting* respectively. Similarly either the traders themselves or the institution can decide when allocations will change. The traders can do this by indicating acceptance of a bid of ask, and the institution can do this by encoding it in a rule, for example that traders whose bid and ask values cross (in other words where a buyer is willing to pay more than a seller is asking for) must trade immediately. We distinguish these cases as *institutional trade-determination* and *non-institutional trade-determination* respectively. Both of these are aspects of clearing.

Whatever form of clearing is used, the effect of the discrete time form of the double auction is that all traders establish their final allocations at the same time. (The final allocation may be the same as the intial allocation of course). The *continuous* or *open-outcry* double auction, in contrast, does not have a specified time for clearing. Instead, in non-institutional trade-determination, buyers and sellers can choose to accept a bid or ask, and then update their allocation, at any point in time. In institutional trade-determination, the institution effectively has to try to clear the auction (which means find a bid and ask that cross) any time a new bid or ask is made by a trader. As Friedman and Rust [16] note, the New York Stock Exchange traded as a call market until the 1860s and its move to continuous trading was driven by growth intrading volume[7], suggesting that it is superior throughput that is the reason for selecting the continuous rather than the discrete-time form of the double auction. However, there can be a cost for this increased throughput—as Gode and Sunder [21] show for their zero-intelligence agents (see below) the expected efficiency of a call market is higher than that of a continuous double auction, and is is easy to verify that this will be the case for other trader behaviors as well.

The rules of the institution thus define the way that the market clears, and the division between markets using different kinds of clearing is the most obvious distinction to be made between double auction variants. However, there are other important distinctions. One of these is what conditions, if any, that the market puts on successive bids and asks. Gode and Sunder [20] give a typical rule when they say that:

> ... any buyer can enter a bid by stating his or her identity, unit price and quantity. The same buyer or other buyers can subsequently raise the bid.

Here unit price is the price that a buyer is prepared to pay for a single good, and this is constrained to increase from bid to bid until a transaction occurs. Such a restriction is called a *bid/ask improvement rule* in [21]. Cliff [6] notes that since this rule is enforced by the New York Stock Exchange, (NYSE) it is often referred to as the NYSE rule (or, as in [48] as "New York rules").

Obviously, in order to be able to ensure that they make a higher bid (or lower ask which is the natural complement of a higher bid), every agent needs to know what other agents bid and ask. However, this kind of knowledge is not

[7]See footnote 24 for more on this.

always available. In a classic open-outcry market (and here the classic market is the commodity trading pit at the old Chicago Board of Trade [16]), every trader is able to hear every bid and ask made (or at least has the potential to do so) [8]. However, less priviledged traders [14] do not have access to current trades. Something similar is the case at the New York Stock Exchange, where *specialists* [25] have complete access to all bids and asks for a type of good (and there is one specialist for each type), while *floor traders* are allowed to check the unmatched bids and asks, so some will periodically have this information while others will only know the value of the highest unmatched bid and lowest unmatched ask ([77] states that this is the most common information provided by a continuous double auction). This latter kind of information is known as the *bid-ask quote*. It is also possible to imagine traders just having information about the price at which the last trade was conducted (which bounds the winning bid and ask rather than giving it precisely) or having no information at all (so any offer has to be based on the trader's private value alone). Taking all of these together we can characterise an auction by the kind (or kinds) of information provided to the traders.

In a continuous double auction it makes sense for traders to get some information about prices, and indeed this kind of feedback is seen as "a major feature" [61] since it is the only way a trader can identify how others value the good and therefore decide how to bid (in other words bids and asks are taken as indications of private values) and as Rust *et al.* [55] point out, this allows auction participants to learn about one another's preferences. Although some call-markets do not issue price quotes—Friedman and Rust [16] point out that this effectively reduces the call market to a kind of sealed-bid auction—it makes sense for them to offer such information for the same reason. The importance of being able to learn is hard to overemphasize. The double auction, whether continuous or periodic (provided there is more than one period):

> ... appears to be too complex a game to yield a clear game-theoretic solution [21][9]

Since it is so hard to predict how other traders will behave in this kind of market, it is not possible to figure out the best way to behave before the auction starts — the only way to determine the best way to behave in a double auction is to learn that behaviour as the auction unfolds.

Another related issue is what happens to unmatched offers. Clearly if a trader changes its bid or ask, then the old one will be discarded, but should anything be done to ones that are (or haven't yet been) updated at the time

[8] This is the kind of market described in [48] where:

> Typically bids are tendered verbally. An auctioneer, upon hearing a bid (offer) writes the bid (offer) and the index of the agent on the chalkboard. The bid (offer) is repeated verbally and the floor is then open for new bids and asks.

[9] This sentiment is echoed by [55], and supported by the fact that all attempts to analyse double auctions that we are aware of, for example [21, 59, 77], only try to analyse call markets and do so under very restrictive assumptions.

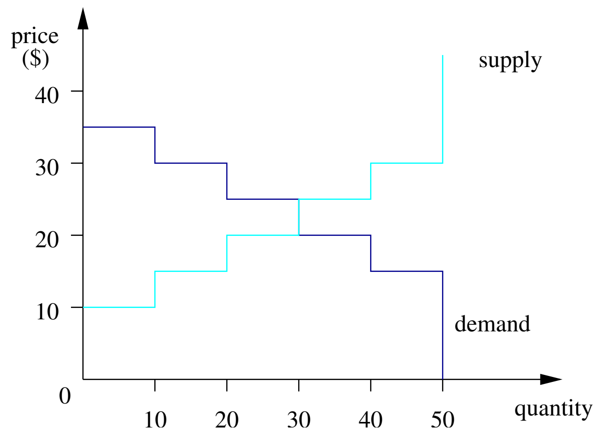

Figure 2: Illustrative supply and demand curves for a double auction.

the auction clears? Unlike the situation with price quotes, there are just two choices. Either the unmatched bids are left open until either they are filled or the auction closes, giving a *persistent shout* auction—a feature that [66] argue is true of real continuous double auctions—or they can be cleared as in [21] and [6].

The first quote from Gode and Sunder [21], above, also illustrates another distinction that can be made between types of double auction. Some are, implicitly or explicitly, *single unit* markets which only allow traders to bid or ask for individual items at a time. Others are *multi-unit* markets which allow traders to bid and ask on several units (and indeed change the number of units between offers). Single unit markets can, and often in experimental work do, include agents that wish to trade multiple units, but only trade a single unit at a time.

Another important distinction between different markets is what mechanism is used to establish the price at which a trade occurs (as opposed to when a trade can occur and if it does when it can, as discussed above). In situations such as that in Figure 1, there is a single price at which trading will be acceptable, \$25, and so this will be the trading price. However, in the market described by Figure 2, there is a *price tunnel* between \$20 and \$25—any price in this range will be acceptable. Price setting mechanisms are required because in general supply and demand curves overlap as in Figure 2. These mechanisms, in the main, make most sense where price-making is institutional. In double auctions with non-institutional price-setting, the price is usually determined by the first offer to be made of the pair of offers that end up being matched — the second offer is considered to be an acceptance of the first.

In call markets, it is typically the case that several buyers and sellers have bids and asks that match when the market clears. Since they are all trading identical goods (at least in all the markets and studies we are considering here this is the case), there are two obvious ways we can decide the trading price for each matched pair—either we can choose one price and apply it to all trades, or we can choose prices for each trade individually. The first is generally termed a *uniform price* [77], the latter a *discriminatory price*. One can easily see arguments for both choices. Uniform price auctions ensure that no trader is discriminated against because of a wayward bid (caused by a misreading of the market for example)[10]. Discriminatory price auctions ensure that the price a trader pays is determined by its bid and that of its trade partner, and supposedly this reflects their amount that they value the good being traded.

Continuous double auctions are, almost by definition, discriminatory-price[11], although it is possible to imagine some kind of institutional price setting that matches traders continuously, but then sets the trade price uniformly at the end of the trading "day", perhaps to the equilibrium price that the market reaches[12]. Discrete-time auctions are typically uniform-price, with transaction price set to the price that gives equilibrium, but one can also imagine a discrete-time auction in which prices are set using a discriminatory mechanism (though there is no obvious reason for doing this).

Whether an auction is uniform or discriminatory priced, it needs to have a mechanism for choosing the price, and there is no standard mechanism for doing this. Instead a number of different mechanisms have been proposed and studied in the literature. [77] studies the application of the mth price mechanism to a uniform-price call market—here the trading price is that of the mth highest bid at clearing, a price that will be less than or equal to the highest matched bid (obviously) and thus greater than or equal to the lowest matched ask. [59] also studies a version of the the call market. In this variation, offers are sealed and so only known to the auctioneer, and the so-called k-double auction mechanism (k-DA) rather than the mth price mechanism is used to set the trade-price between the lowest matched bid and highest matched ask[13]. For the k-double auction mechanism we choose a value k between 0 and 1, and then set the price to be:

$$kb + (1 - k)a$$

where a is the ask, b is the bid, $b > a$, and $[a, b]$ is the interval in which a price can be selected. One could either apply this mechanism in a uniform-price auction (where the interval is taken across all matching bids and asks) or in a discriminatory auction, and it is also possible to use it in a continuous double auction. [77] call the parameter κ rather than k when it is used in a discriminatory rather than a uniform-price rule.

[10]Though its bid will typically help to set the trading price, through a mechanism like picking the mth price.

[11]Indeed Madhavan [38] distinguishes continuous and discrete-time auctions precisely by the fact that continuous auctions use discriminatory-pricing and discrete-time auctions use uniform-pricing.

[12]We have recently [47] been looking at a form of auction which, though discriminatory price, tries to learn the equilibrium price, and in doing so sets the transaction price in a way that is affected by previously matched offers as well as the ones involved in the transaction.

[13]The lowest matched bid is the lowest bid to exceed at least one ask, and the lowest matched ask is the lowest ask to exceed at least one bid. Thus these two offers are the ones on the boundary of trading, and a price that is set between them will be acceptable to every trader with a matched offer.

The version of the k-double auction obtained when $k = 1$, and the price is determined by the lowest matched bid, is called the *buyer's bid double auction* (BBDA) [58]. A variation on this, the *modified* buyer's bid double auction (the MBBDA [32], or MDA [59]), where the price is determined not by the lowest matched bid, but by "the lowest bid to the left of the intersection of the supply and demand (curves)"[14] [32], has also been studied [32, 59]. Note that both the buyer's bid and modified buyers bid double auctions involve sealed bids, in contrast to the open bidding of a typical call market, and since the buyer's bid double auction is that obtained by running a k-DA with $k = 1$, this mechanism is also known as the 1-DA. we should note that although the k-DA mechanism, as introduced, works with sealed bids, there seems to be no reason why it should not work perfectly well for an open-outcry market.

A k or κ double auction treats all trades of a given type alike — the parameter sets the fraction of profit that goes to buyer and seller, and all buyers and sellers get that fraction. The buyers bid double auction and the modified buyers bid double auction behave in the same kind of way. The mechanism described above, where the price of a trade is the price of the bid or ask that was made first, and so it "on the table" when the match occurs, will clearly allocate all of the profit in the transaction to the trader who made the deal possible. This trader can be either buyer or seller, so this mechanism does not reduce to any version of the k or κ double auction.

Another aspect that needs to be considered is the timing of the offers made. In much of the theoretical work that has been carried out, it is assumed that offers are made simultaneously, since this simplifies the analysis), creating a kind of double sealed-bid auction. Real markets, in contrast, deal with offers that arrive over time, allowing traders to make inferences from the offers made by others. Indeed, in many of these markets market quotes play an important role in providing information that allows traders to decide on what prices to offer.

Now, so far we have assumed that all trading takes place between peers—between agents that either buy or sell—and that the auctioneer organises the market but does not trade. However, there are markets in which the role of auctioneer is taken by a *market maker* who also makes trades. [38] distinguishes between *quote-driven* and *order-driven* markets. The latter are the kinds of auction that we have discussed up to now, markets in which traders submit orders and these are either matched with existing orders, or wait on an order book until a matching order is received by the auctioneer. Quote-driven systems, like those that operate on NASDAQ and the London Stock Exchange (LSE), are *continuous dealer markets* in which traders who wish to execute an order can trade directly with the market maker — it is part of the market maker's job to provide liquidity in this way, always buying or selling goods on demand. While the continuous dealer markets on the LSE, NASDAQ and the specialist part of the NYSE have only one market maker per type of stock (the specialist for that stock), other markets, for example the Chicago Board of Options Exchange

[14] In other words the bid that is the next highest bid after the lowest matched bid.

(CBOE), operate as *multiple dealer markets* [36] where several market makers operate and traders choose which one to deal with. From the traders' point of view, the advantage of a quote-driven system is that they can always find out what price they will trade at before submitting an order (the current quote price), whereas in an order-driven system prices need not be determined until a match is made.

So, to summarise, we can distinguish between double auctions by answering the following questions:

- Is the market a quote-driven or an order-driven (auction) market?
- If the market is quote-driven, is it a specialist market or a multiple dealer market?
- Is the market one-shot or repeated?
- Is it periodic or continuous?
- How do offers arrive over time?
- Are prices set by an institution or by the traders?
- Is the decision to trade taken by the traders or by the institution?
- What conditions, if any, are placed on successive bids and asks (is there a bid/ask improvement rule)?
- What information about other traders offers does each trader have?
- What happens to unmatched bids and asks when a match occurs (are they deleted, or do they persist)?
- Do the buyers want to buy one or many goods? Do sellers want to sell one or many goods?
- How many goods are bid (or asked) for at a given instant by one of the traders (one or many)?
- Are trades priced using discriminatory or uniform pricing?
- How are the uniform or discriminatory prices determined?

These answers to these questions will identify the point that a particular auction occupies in the space of all possible instantiations of the parameters to which these questions refer. The set of parameters overlaps with, and is not (completely) subsumed by the parameterization in [78] (as can be seen by comapring the above with Table VIII [78][p. 335]).

We will finish this section with the description of a particular double auction market taken from [37, pp. 186–92] via Smith [64, p. 945] which neatly specifies the values of a number of the parameters described above:

> After the market opens an auction for a num

> it begins with the announcement of a price bid by any buyer or a price offer by any seller. Any subsequent bid (offer) must be at a higher (lower) price to be admissable. Once a bid (offer) has been made public, that is, is "standing", it cannot be withdrawn. A binding contract occurs when any buyer (seller) accepts the offer (bid) of any seller (buyer). The auction ends with a contract. If the standing bid (offer) was not part of the contract (for example, if a buyer other than the one with the standing bid accepted the standing offer), the maker of that bid (offer) is no longer bound to it unless the bid (offer) is now re-entered. Following a contract a new auction begins when a new price bid (offer) is announced. The new bid (offer) may be at any level, and may involve "signalling". This process continues until the market "day" comes to an end.

As Smith notes [64], there are a number of variations that fit within this overall description (some of which are studied in [65]), and it is close other descriptions that one can find in the literature (for example that in [48][p 1493]).

2.2 How to measure performance of auctions and traders

The reason behind much of the work on auctions has been to try and quantify the performance of either auctions or traders. Here we discuss some of the measures that are most commonly used to do this.

We start by considering the profit of individual traders. If the ith transaction matches a buyer with private value v_{b_i} and a seller with private value v_{s_i}, setting a price p_i, then:

$$pr_{b_i} = v_{b_i} - p_i$$

is the profit that the buyer makes from the trade, and:

$$pr_{s_i} = p_i - v_{s_i}$$

is the profit that the seller makes. If a trader does not take part in a transaction, then its profit is 0.

Economists have traditionally been more interested in measures of performance of an entire market than in the profit made by individuals in that market. Given the private values of all of the agents in a market, we can calculate not only the profit that each individual makes, but also how the total profit, or *surplus* actually made, S_a compares to the maximum profit that theoretically could be made S_t. This measure is the *allocative efficiency*, e usually expressed as a percentage.

$$e = \frac{S_a}{S_t}$$

A market that is 100% efficient extracts all the possible profit from the traders that operate within it.

The actual surplus is easy to compute from the profit of the individual traders:

$$S_a = \sum_{i=1}^{n} v_{b_i} - p_i + \sum_{j=1}^{n} p_j - v_{s_j}$$

for all n transactions. The theoretical surplus can be computed by establishing the *equilibrium price*, p_0. The equilibrium price is defined as the price at which supply and demand are equal. In markets like those in Figure 1, as discussed above, this is a unique price. In markets like those in Figure 2, this only defines a range of prices, and in such cases we take it to be the midpoint of the price range (the price that would be set by a $k = 0.5$ auction).

Given the equilibrium price, the theoretical surplus is computed as the sum of the profits that would be made by buyers and sellers were all transactions to have taken place at the equilibrium price p_0 — the *equilibrium profits* epr of the traders. In other words:

$$epr = epr_b + epr_s$$

where epr_b and epr_s are the equilibrium profits of all buyers and all sellers respectively. These can be calculated as

$$epr_b = \sum_{i=1}^{n} epr_{b_i}$$

where the equilibrium profit for each buyer is:

$$epr_{b_i} = v_{b_i} - p_0$$

and similarly,

$$epr_s = \sum_{j=1}^{m} epr_{b_i}$$

where

$$epr_{b_i} = v_{b_i} - p_0$$

for each seller. Thus, overall, allocative efficiency is given by:

$$e = \frac{\left(\sum_{i=1}^{n} v_{b_i} - p_i\right) + \left(\sum_{j=1}^{m} p_j - v_{s_j}\right)}{\left(\sum_{i=1}^{n} v_{b_i} - p_0\right) + \left(\sum_{j=1}^{m} p_0 - v_{s_j}\right)} \tag{1}$$

Since allocative efficiency will not detect whether a market is skewed towards the buyer or seller, alternative measures have been proposed, including the notion of buyer and seller *market-power* [46]. The idea of the market-power measures is to determine how much profit the buyers or sellers actually obtained as opposed to how much they would obtain were the market trading at the equilibrium price. In other words, how much have buyers or sellers managed to tip things in their favor. Buyer market power mp_b is calculated as:

$$mp_b = \frac{1}{n}\sum_{i=1}^{n}\frac{pr_{b_i} - epr_b}{epr_b}$$

where i ranges across all n goods sold, pr_{b_i} is as computed above, and epr_b is the equilibrium profit of the buyers (as above). Seller market power mp_s and equilibrium profits epr_{s_i} and epr_s are calculated similarly . Another measure which aims to quantify the same aspect of trader performance, that is the relative profitability of the trader, is the *profit dispersion* measure suggested by Gode and Sunder [20]. This is the root mean squared difference between the actual profit and the equilibrium profit of individual traders. For m traders, profit dispersion pd this comes to:

$$pd = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(pr_i - epr_i)^2}$$

where epr_i is the equilibrium profit of that trader, either epr_{b_i} or epr_{s_i}, while pr_i is the actual profit, either pr_{b_i} or pr_{s_i}. Note that here i ranges across all traders rather than just across buyers or sellers.

The main difference between these two measures is the following. In market power calculations, profits above equilibrum for one trader will balance profits below equilibrium for another. In other words an increase in market power for buyers will be matched by a decrease in market power for sellers. In contrast, any deviation from equilibrium will cause an increase in profit dispersion.

Measuring market power gives a finer-grained view of what is happening in an auction than just looking at allocative efficiency, but it is still an aggregate measure across all buyers or sellers. As a result, there are aspects of market behaviour that it cannot detect. In a market with two types of buyer, say, one that gains high profits and one that makes big losses, the market power may come close to 1 even though it would be far from 1 (both positively and negatively) for the two sub-populations. As a result, it can be useful to measure the profits made by individual traders, or groups of traders, either in absolute terms, as a proportion of total profit obtained (as in [66]), or in a way similar to market power.

Profit is an important measure for markets, but it is also interesting to consider how close real markets come to trading at the point where theory says they will trade, that is at the equilibrium price. Again this is important because just considering the profit (and in particular the allocative efficiency) can obscure some strange behaviour (ZI-C agents (see below) for example can trade with high allocative efficiency but trade far from the theoretical equilibrium price). To measure how close a market is to the theoretical equilibrium, Smith [61] introduced the idea of the *coefficient of convergence*[15], α such that:

$$\alpha = \frac{100\sigma_0}{p_0}$$

where p_0 is the equilibrium price and σ_0 is the standard deviation of trade prices around the equilibrium. This can be reformulated as:

$$\alpha = \frac{100}{p_0}\sqrt{\frac{1}{n}\sum_{i=1}^{n}(p_i - p_0)^2}$$

where p_i is the trade price of transaction i.

While α gives us a way of determining, on average, how close a market trades to the theoretical equilibrium, it will conflate two aspects of "closeness" — the speed with which prices converge (since prices typically "scallop" [10] as they converge on the equilibrium price), and the deviation from the equilibrium price at which transaction prices eventually settle. There does not seem to be a measure that captures this, so that, for example, Preist and van Tol [49] do not have a quantitative measure with which to back up their claim that their trading algorithm converges faster than ZIP, and all that they can do is to plot transaction prices over time for a given run.

To summarize, when comparing auctions and the bidding strategies used in them, we are typically interested in answering questions such as:

- What is the allocative efficiency of the auction?
- What is the buyer and seller market power?
- What is the profit dispersion?
- What is the profit that individual traders (or types of traders) a ccrue during the auction?
- What is the coefficient of convergence for the auction?

2.3 Categorising some double auctions

In this section we use the attributes identified above to categorise some of the double auctions discussed in the literature. Two things should be noted about this. First, we don't aim to give a comprehensive description of markets in this section, we are more interested in showing the range of different markets that one can find. Indeed, we only consider a handful of different markets described by Smith [61], Rust *et al.* [54, 55], Gode and Sunder [20, 21], Cliff [6, 9, 7, 8], Preist and Van Tol [49], and Das, Tesauro *et al.* [10, 66, 67]. Second, we are focused on the structure of the markets not on the work carried out on those markets, so the description of a number of the works in the literature is

[15] Frequently referred to as "Smith's alpha", as in [49].

rather incomplete — we provide a more extensive description of this work in Sections 4 and 5. Note that the development of electronic markets, and many of the markets that have been studied are electronic, rather blurs the distinction between real and simulated markets and both kinds are convered here. As a result, some of the distinctions we identify are artifacts of the way that the simulations are performed — aspects that will be important to anyone who is interested, as we are, in replicating results.

We start with the market described in Smith's classic study of competitive market behaviour [61][16] in which the traders were all human. In the experiments, each trader was given a specific limit price of a unit of a "fictitious commodity" (all of these prices are specified in the paper) and all traders were interested in trading a single unit of the commodity (at least in the first eight experiments described in [61]). Buyers were constrained to purchase a unit for less than their limit price, and sellers to accept a price no lower than their limit price.

The market mechanism specifies that traders are allowed to make offers at any time. Buyers make offers by calling out the price at which they are willing to buy, sellers make offers by calling out the price at which they wish to sell. A buyer may accept an existing offer to sell, and, symmetrically, a seller may accept an offer to by, creating a contract. When this happens both traders take no further part in the market until the end of the trading period. Offers seem to be persistent — there is nothing in [61] to say they are erased by a trade as they are in some experimental work — though presumably traders forget old offers, and all traders can hear all offers. When traders are described as "accepting an offer", it indicates that the trade takes place at the price of the initial offer.

Smith ran the markets for several periods (the number varied from experiment to experiment, but was typically around 4 or 5) with each trader having their supply of, or demand for, goods replenished between periods so that the market is one in a steady state (except for experimentally applied supply or demand shocks). The final aspect of [61] that we will note is that in later experiments (9A, 9B and 10), traders are allowed to trade two units at their given limit price, showing that:

> doubling volume in a given period is comparable to running two trading periods at the same volume.

The simulated auction used in the Santa Fe Double Auction tournament [54, 55] comes close to being an implementation of Smith's experimental markets. At the finest level of detail, the Santa Fe auction is structured as a series of "steps", where one step allows bids and asks to be posted and in the subsequent step, if bids and asks cross, the traders with the highest bid and lowest ask decide whether to accept the other's offer. A trade occurs if either makes that decision. The trading period is a fixed number of these steps (the number is known to the

[16] This paper, published in 1962 reports experiments "performed over a six-year period beginning in 1955. They are part of a continuing study . . . ", and later experiments (some of which are described below) deal with more complex scenarios.

traders), and a number of these periods make a "round", of which there may be many. The decision to trade is made by individuals, though only the matched individuals are allowed to trade[17], and it is not clear how the transaction price is set, though a bid/ask improvement rule does seem to be enforced. In addition it seems (according to [66]) that unmatched bids are cleared when transactions occur, that all traders have access to all the offers made by all the other traders, and that only one good is bid for at a time.

Gode and Sunder's work [20] on zero-intelligence agents concerns repeated auctions (and each auction but one in their study were run for six trading periods), the auction was continuous, and only a single unit was traded at a time. Trade determination is institutional, in the sense that a binding transaction occurs as soon as a bid and ask cross, and price determination also seems to be institutional — as described above, the price is determined by whichever of the bid or ask is already "on the table" when matching occurs. In addition, unmatched offers are cleared when a transaction occurs, and only one good is bid for at a time. All of these factors are clear from the description of the experiments.

The remaining factors are less clear from the paper. It seems that though Gode and Sunder consider a bid/ask improvement rule to be an essential part of a double auction in their theoretical work [21], there seems to be no mechanism for it in the automated traders they discuss (indeed since these are described as having "no memory" it would seem impossible for such a mechanism to be implemented), and it is not clear whether such a rule is applied to the human traders in that study either. Furthermore, it is not clear what information traders have. From the description of the experiment with human traders, it seems inevitable that the human traders knew all the offers that had been made. There is nothing to say whether the automated traders had this information, but if they did, then the mechanism they used for choosing bids and asks did not make use of it.

Somewhat clearer is the setup Gode and Sunder consider in [21]. Here they analyse two kinds of double auction, a *synchronized* auction, which falls somewhere between the classic discrete time auction and a continuous double auction, and a *continuously clearing* auction (called *continuous* in [22]), which is much closer to a classic continuous auction. In the synchronized auction, all traders are simultaneously polled for offers — so that traders do not know other traders' offers when they make their own — and if the highest bid and lowest ask cross, then a binding transaction takes place. (Note that it is not clear what happens if several traders tie for the highest bid or lowest ask.) If the highest bid and lowest ask don't cross, then traders are polled again and submit new bids under a bid/ask improvement rule. As soon as there is a transaction, all unaccepted bids and asks are cancelled, and the process is repeated until the pre-specified end of the period. If all matching bids and asks were cleared after a poll, then this auction would be close to come implmentations of discrete time auctions.

[17] As the authors point out, these are the same rules as adopted by the AURORA system developed by the Chicago Board of Trade.

The continuously clearing auction is intended to be closer to a real auction where

> the timing of ... bids and asks is determined by the free will of individual traders, and a transaction is completed as soon as a bid and an ask cross, without waiting for the remaining traders in the market to submit their bids/asks.

To simulate this, Gode and Sunder model traders making offers by randomly sampling traders, without replacement, to make an offer, so every trader gets to make an offer before any trader makes a second offer (so there is still some synchronisation). As soon as a bid and ask cross a transaction is executed, and all unaccepted offers are cancelled, and if the first round of bidding doesn't result in a transaction, further rounds are carried out with the order in which offers are solicited being randomly generated each round.

In the continuous auction, a trader is randomly picked to make an offer, that may update the highest bid and lowest ask. In [22], the authors distinguish between the trading being picked with and without *replacement* (in the usual probabilistic sense). When traders are picked without replacement, no trader is able to make a second offer before all other traders have made their first offer, and there a clear rounds of offer-making. With replacement, a given trader may be allowed to make an arbitrary number of offers before any other trader gets to make an offer (though this is, of course, unlikely), and there are no clear rounds. [22] also introduces an auction that Gode and Sunder call *semi-continuous* (or *hybrid*. In this form of auction, traders are polled for offers as in the synchronized auction. Then offers are randomly selected from this pool, as in the continuous auction, and the current highest bid and lowest asks tested to see if they cross. This selection is carried out without replacement.

In Cliff's work on zero-intelligence plus (ZIP) agents [6, 7, 8, 9][18] the auction is repeated and continuous, and there are no persistent shouts—all bids and asks are cleared when a trade takes place. Furthermore, as described in [6, pp. 100–111], the way that the auction is simulated is quite specific. At each simulation round or "trading session", one agent from those currently able to trade is randomly selected to make an offer for a single unit. All other agents are made aware of this offer. If the agent that made the offer was a buyer, then all sellers are polled to see if they want to respond by accepting the offer. If one seller does want to take advantage of the offer, then trading takes place at the price suggested by the first agent, just as in [20]). If several agents respond, one is chosen at random to complete the trade. If the agent originally picked ot offer was a seller, then the symmetric process follows. Rounds are repeated either until either no more buyers or no more sellers can trade, or until some

[18] These four papers overlap to a large degree. Without knowing the history of their writing, it seems that [6] was written first as a complete summary of the original work (including a listing of the code used to run the simulations), [9] was a slimmed down version without the background on market-based control and background economics, and then [7] and [8] were then extracted out as conference papers (the former introducing ZIP traders and the latter identifying problems with ZI-C traders).

limited number of rounds has passed without a trade. Note that, as in [20], offers are cancelled by a trade.

From this description, we can see that the market is synchronised, with each round producing at most one trade. However, the auction that is created is neither a continuous double auction nor the kind of auction described in any of the work above. It comes closest to being a continuous double auction, but differs since at each round some traders are not able to make an offer — those traders of the same type, buyer or seller, as the agent originally picked to make an offer. The fact that all agents see the offer minimises the effect of this, but the approach makes it impossible for one trader, a buyer say, to scoop the offer made by the randomly picked agent at the start of the round. For example, if that randomly picked agent is a buyer, no other buyer can step in with a higher bid before the sellers respond[19].

Preist and van Tol [49] build directly on Cliff's work, examining the performance of a similar kind of trader, and so use much the same experimental structure. Time is again divided into "rounds", with all agents having to make an opening offer at the first round, and any agent being able to update its offer at a later round (unless of course it has traded), simulating a continuous auction. Offers persist unless updated. Offers that cross are cleared just as in Smith's work, and Preist and van Tol argue both that their market is close to Smith's — differing in that automated traders will not forget past offers — and is more flexible than Cliff's.

All the markets described so far have been designed for ease of carrying out experiments with automated traders rather than for realism. As a result, that of [61] is the only one that has been a true continuous double auction of the kind one finds operating in real markets. Das, Tesauro *et al.* addressed this issue [10, 34] explicitly, going to great lengths to ensure that offers could be made at any time and in real time, also allowing for persistent offers. The market runs over several trading periods (15–16 periods, each of 3 minutes duration), with private values being held constant for some periods and then changed. Other market parameters were as follows: a single unit was traded at a time, players had several units to trade in each period (8–14) and units/money were replenished at the end of each period, trades occurred as soon as bids and asks crossed, the transaction price was the price of the first of the two offers, a bid/ask improvement rule (again called the "NYSE rule") was in force.

Further work by these authors [66, 67] used a similar setup[20], with the exception that real-time simulation was replaced by a discrete-time simulation. The discrete-time model generalises that of previous work, picking agents to be active in a given time-step with some probability α_i, so that at any point there

[19] Imagine we have three agents that have still to trade, buyer1 with private value 20, buyer2 with private value 30, and seller1 with private value 10. If buyer1 is picked to offer, and offers 15, then seller1 may well accept before buyer2 can make the offer of 25, say, preventing an outcome with greater surplus.

[20] Though the description of the setup in [67] suggests that experiments were carried out under a broader range of conditions, without persistent offers for example, the results given were under broadly the same conditions as above, with the exception of the discrete event simulation described here.

can, in theory, be any number of agents active at a given time (when there are two or more, their offers are handled in a random order).

What we can see from this cross-section of work is that research has aimed to be representative rather than strictly reproducing markets that really occur — for example the fact that many approaches enforce some synchronization of offers — but has covered a range of subtly different markets, something that makes it hard to directly compare the results of the work. This problem is exacebated, as we will see in Section 5 when we look at experimental work using these markets, because the performance of the traders is determined in different ways as well. However, before we look at the kinds of results that have been obtained with these different setups, we will look at a specific market in some detail.

3 The New York Stock Exchange

Have described in the abstract some of the different kinds of double-sided market, it is interesting to consider a concrete example of such a market in the form of the New York Stock Exchange (NYSE). The description given is drawn from that in [26]. As hinted at above, the NYSE is rather more complex than anything we have described so far.

Physically the NYSE is split into the *upstairs* and *downstairs* markets, where the downstairs market is the main trading floor (we will come back to the upstairs market later). The main trading floor is split into a number of rooms[21], throughout which are spread seventeen trading posts. It is at the trading post for a specific stock that the specialist for that stock works, and it is at the post that the *trading crowd* for that stock will gather. The trading crowd is made up of *floor brokers*. Floor brokers are allowed to trade in any stock and to operate at any trading post, but typcially concentrate on a few posts and trade only in the stocks traded at those posts.

Trading on the NYSE takes place between 9.30am and 4pm EST Monday to Friday. After opening, trading in each stock takes the form of a continuous auction[22], with each specialist receiving all offers for the stock in which they deal, and maintaining the order book that records such offers. The specialist supervises the trading process, makes matches between buyers and sellers, and is responsible for crowd control amongst the trading crowd. As described above, the specialist is required to trade when necessary, and can also trade when it is advantageous to do so, but as [40] points out, the volume of trade involving specialists is now a minority, less than 20% by 2000. The vast bulk of trades are direct between brokers.

When the exchange is not open for trading, orders continue to arrive elec-

[21] These rooms are called "The Garage", the "Main Room", the "Blue Room" and the "Expanded Blue Room" [26].

[22] Though, as [38] puts it, "for some thickly traded stocks, the market has been described as a continuous double auction" which suggests that the market in most stocks is too thin to be regarded as a true CDA.

tronically, and are stored in the Opening Automated Report Service (OARS), which matches bids and asks. In addition, prior to opening, floor brokers who want their orders to figure in the opening may pass orders to the relevant specialist, and these are also entered into OARS[23]. When the market opens, the specialist sets a single price at which initial trades must be executed, and this price may be some way from the price at which the market would theoretically clear given the public bids and asks [40].

Thus we can think of the NYSE, at a first approximation, as using a call market to establish opening prices, and then trading as a continuous double auction — the approximation relating to the fact that this description ignores the role of specialists — and this is often the way that the NYSE is presented. These aspects of the NYSE are common in other stock exchanges. [38] reports that the open-outcry system for stocks on the Paris Bourse, the Swiss Option and Financial Futures Exchange (SOFFEX), the Frankfurt Stock Exchange, the Computer Assisted Trading System of the Toronto Stock Exchange and the Computer Assisted Routing and Execution System of the Tokyo Stock Exchange all operate as continuous double auctions, while discrete-time auctions are used to open the NYSE and the Tokyo Stock Exchange. Other stock markets operate purely as discrete-time systems, the Arizona Stock Exchange, which operates as a Wunsch auction [44].

The role of the specialists in the NYSE is studied in detail by [25], who both expand on their role in the institution, and examine empirical evidence about their activity in the market (which, for example, suggest that they are typically good short-term traders but not particularly good long term traders). Specialists are required to maintain price continuity—which basically limits successive price changes to one eighth of a dollar—and to act to stabilize the market, buying when the price drops and selling when it rises. These requirements expose the specialist to sudden markets shocks, but this disadvantage is offset by the fact that the specialist has complete knowledge of the order book—the specialist has to allow floor traders to view the order book, but the book is not widely displayed so the specialist has a significant advantage in terms of information [40]—and gets to see incoming electronic orders before these are posted in the order book.

The upstairs market [39] is an informal market that aims to deal with large blocks of shares (where a "large block" is well over 10,000 shares) by negotiating a price between potential buyers and sellers. The idea seems to be that it may be hard for the trading floor to absorb such blocks, though as [26] points out, only around 15% of total trade is arranged through the upstairs market.

There is one final quirk of the NYSE that is worth noting. The exchange itself is a not-for-profit organization, founded originally in 1792 [24], that is owned by

[23]As described in [40], brokers may make orders for the open without being physically present, and may choose to participate in a certain percentage of the specialist's opening trade instead of placing a standard limit order.

[24]According to [2, page 310], the recent authorative history of New York City, trading in public stocks began two or three years prior to 1792, with most stock being sold off in a form of English auction. Sales were nominally public, but most stock was purchased by "jobbers"

those who trade in it, so-called *seatholders*[25]. Seats are limited—there were originally 1,060 seats, and this number was expanded to 1,375 in 1932 and 1,366 in 1953—and seatholders are the only people allowed to operate on the floor of the exchange, either as specialists, floor traders or brokers as they wish[26]. Since 1978, seats may be leased, and this is now relatively common, but seats may also be bought and sold (seats were first traded in 1869), and this trading itself takes place through an ongoing auction market run by the NYSE. This market is theoretically continuous, and the bid and ask quotes are posted on the trading floor, but seats are thinly traded.

---

or "dealers", who bought for themselves, and "brokers", who bought for clients. Auctions were held several times a week — [2] give the example of the auction run by the partnership of Bleeker and Pintard, which was held at the Merchant's Coffee house, three times a week at 1pm and three times a week at 7pm — and between auctions deals were done direct between the various brokers and jobbers in the "curb market". In the summer of 1791 these extra-auction dealings were held under a sycamore tree on Wall Street, in September of that year, as the volume of trade increased, the auctioneers, jobbers and brokers agreed a set of rules for stock trading, and in February 1792 the first "Stock Exchange Office" was opened on Wall Street.

Shortly after this, in March 1792, the fledging stock exchange, which was still operating as an English auction, collapsed, largely because the market had been manipulated on a massive scale — partly by the auctioneers, partly through rumor-mongering by speculators — and stock was greatly over-valued. During the fallout from the collapse, a group of dealers and brokers agreed not to trade stocks by English auction, and in May a further group made the"Buttonwood Agreement" ("buttonwood" being a New England name for a sycamore) setting a minimum commission rate and stating that preference in trading be given to other signers [2, page 311]. By May of 1793, the market was being held upstairs at the Tontine Coffee House, at the intersection of Wall Street and Water Street, where it remained until 1827 (along with the curb market, which continued outside).

By 1817, [2, page 445], the volume of trade was such that the Buttonwood Agreement was revoked and the stock market was reorganised under the auspices of the new New York Stock and Exchange Board. During this time trading was still carried out at the Tontine Coffee House, but, in 1827, the exchange moved into the newly built Merchant's Exchange building. This building burnt to the ground in the great fire of 1835 which devastated lower Manhattan (also claiming the Coffee House) and the Stock Exchange had to find alternative accomodation until the Merchant's Exchange was rebuilt several years later. The Stock Exchange then remained in the Merchants Exchange, side-by-side with out exchanges like the New York Produce Exchange until it finally moved into its own building in 1865 [2, page 877].

Before that move, however, had come another major change in the rules of the Exchange [2, page 899]. The massive growth in general trade in New York during the Civil War — including much trade with the Confederacy, not all of it legal — and the concommitant growth in stock trading meant that by 1864 the old methods of trading on the Stock Exchange could not cope with the volume of stocks being traded. As a result:

> A breakaway group of brokers established the Long Room on Broad Street where stocks could be traded continuously and hundreds of transactions could take place simultaneously. In addition "curb brokers" (also known as "guttersnipes") traded in the streets from after the New York Stock Exchange's five P.M. closing until the sun went down. One businessman opened an Evening Exchange, providing the only place in the world where securities could be traded twenty-four hours a day. [2, page 899]

This seems to have been the start of trading on the NYSE as it is carried out today.

[25] Up to 1871, members sat in assigned seats during the roll call of stocks.

[26] There are other roles, as described in [33] but these need not concern us here.

4 Theoretical work on double auctions

This section surveys some of the theoretical results that are pertinent to the study of double auctions. We start by considering work on auctions markets as a whole, before looking at work analysing the behaviour of traders within markets.

4.1 Market behaviour

Historically, auction theory considered single-sided auctions before double auctions, and looked at clearing house auctions before continuous double auctions (even now there is little written about the theory of CDAs). This is doubtless due to the relative complexity of the auctions — single sided auctions are, by and large, simpler than double auctions, and among double auctions, clearing houses are simpler to analyze than CDAs — and the fact that it is natural to analyze the simpler cases before the more complex cases. As a result of this development, there is little theory that can directly help us in trying to understand the functioning of CH and CDA markets, and some of the work that can help us does so indirectly, using results that were orginally derived for single-sided auctions.

Much of the auction theory literature hinges on the existence of an equilibrium in which "buyers' bidding strategies are monotonic in their types" [42], meaning, broadly speaking, that the price a buyer is prepared to bid is a monotonic function of their private value for the good in question[27]. The existence of such an equilibrium can be easily obtained if private values are distributed smoothly, independently, and symmetrically, and this is helpful in the anlaysis of cases which have this characteristic. However, there are plenty of cases which do not have this characteristic, and Maskin and Riley [42] addressed some of these. In particular, they showed that monotonic equilibria exist, under some conditions on preferences, for all cases of independent private values and, under additional conditions on preferences, when private values are affiliated [45][28] provided that tying bids are broken using a Vickrey action [68]. However, Maskin and Riley only considered the case of single-sided auctions. Double-sided auctions are more difficult to analyse.

The first such analysis for a double auction was that of Chatterjee and Samuelson [5], in the paper in which they introduced the idea of the k-double auction[29] albeit only for the two trader case. In this initial paper, Chatterjee and Samuelson show that there is an equilibrium solution, assuming independent private values.

Considerable work has since been carried out extending this result. First, Williams showed the existence of equilibria in the buyer's bid double auction [74,

[27]This statement is accurate enough for our purposes. For a precise characterisation, see, for example [42]

[28]Which, broadly speaking, means that one trader having a higher private value for a good makes it more likely that another will.

[29]Though not under this name — they refer to the price setting rule as a "bargaining rule".

73] — this is an easier auction to analyse since the dominant strategy for sellers is to bid their true value, thus fixing one side of the auction and, as [56] points out, ensuring that the market has a unique equilibrium[30]. The same authors subsequently showed the existence of equilibria in the many-trader version of the k-double auction [59], at the same time suggesting that the modified BBDA has no equilibrium. This work was followed by Jackson and Swinkels [29, 30], who showed the existence of equilibria, though not monotonic equilibria, under a wide range of conditions. Next, Reny and Perry [51] showed that monotonic equilibria existed if offers are restricted to discrete values, and Fudenberg *et al.* [17] showed that this result could be extended to continuous values (which [30] argues is "a very useful approximation ... allowing one ... to use calculus to characterise equilibria") provided that the auction was large. Finally, Kadan [31] showed that an increasing equilibrium exists for just two traders with affiliated values.

In addition to identifying whether auctions have equilibria, it is interesting to ask whether such auctions are *efficient.* we have already met the idea of allocative efficiency (1), and we can say that an auction is efficient if its allocative efficiency is 100%. In other words, an auction is efficient if the result of the auction maximises the profit obtained by the traders. This notion of efficiency coincides with the idea of maximising *social welfare*, ensuring that the goods sold by the auction are obtained by the buyers who value them most, and, symmetrically, are unloaded by the sellers who value them least.

Wilson [75] analyzed what he described as a sealed bid double auction, a clearing house in which traders are not aware of the offers made by other traders, and which we can describe as a multi-trader version of the k-double auction introduced in [5]. ([75] was the first paper to consider this multilateral k-double auction.) The main result of this paper is that the auction is *incentive efficient*, meaning that there is no mechanism that traders will prefer (there is no incentive for traders to prefer another mechanism). This result hinges on there being enough traders in the market, which means, in practice, that the incentive efficiency holds as the number of traders increases to ∞, and the market will satisfy the notion of *interim efficiency* suggested by [27][31].

In a more general setting than we are considering, Hurwicz [28] has shown that it is not possible to construct a mechanism that gives results that are individually rational, Pareto optimal and *strategy proof* — a strategy proof mechanism is for which no individual trader can benefit by not truthfully reporting their private value for the good in question. This means that always has to take into account the fact that traders can benefit from not reporting their valuations for goods in a truthful manner, and, at first sight, this result seems to be a very negative one from the perspective of designing auctions. However,

[30]Note that all results for the BBDA, the 1-DA, are symmetric with those for the k-double auction in which the transaction price is determined by the price offered by the highest asking seller that trades, the 0-DA [56].

[31]This means,broadly speaking, that there is no other mechanism which is preferred by any individual in the market when those individuals do not know the private values of all other individuals, but do know their own private value.

its impact is limited, and researchers have worked to identify the bounds of this limitation. For example, Roberts and Postlewaite [52] show that, as the number of individuals tends to infinity, the gain any trader can make by not being truthful falls to zero. If the market is big enough, there is no advantage in not telling the truth.

A related result is given by Satterthwaite and Williams [58] for the buyers bid double auction[32]. They consider Bayesian Nash equilibria of this market, assuming sellers will ask truthfully (since their offers do not affect the price they pay, they have no incentive to deviate from truthtelling), and show that any equilibrium strategy for buyers involves them bidding in such a way that the difference between their bid and their private value is inversely proportional to the number of traders. Since truthful bidding in the BBDA leads to efficiency, this result shows that the auction converges to efficiency as the number of bidders increases. In [56], Rustichini *et al.* show that the result about the deviation from efficiency also holds for the k-double auction, providing a complementary result to Wilson's [75] proof of convergence (the result also appears in [59]). Thus the k-double auction is no worse than the BBDA in terms of the speed with which it converges to being strategy-proof as the number of trades increases, and numerical results provided in [56] indicate that the inefficiency in a k-DA with six buyers and six sellers should be less than 1%. Such a market is effectively efficient. Subsequently Satterthwaite and Williams [60] have shown that the worst-case speed of convergence of the k-double auction is as high as is possible for any auction.

Note that all this work on convergence to efficiency assumes that traders have independent private values. As [56] point out, it would be better to be able to obtain results for the case where private values are not independent — as we have already seen, we do not yet have results about convergence to efficiency under such conditions.

McAfee [43] has extended this line work by suggesting a mechanism, the so-called "dominant strategy auction", in which every trader has a dominant strategy of telling the truth, *whatever* the number of traders. This mechanism is a variation of the clearing house, with all offers being made before the transaction price is set. The price is determined not, as it is in the k-double auction, by the lowest matched bid and the highest matched ask, but by the next-to-lowest matched bid (the next bid up) and the next-to-lowest matched ask (the next ask down). The transaction price is set midway between these two values, and for some sets of offers the market will clear just as a CH, while for others, those for which the transaction price is not beween the lowest matched bid and the highest matched ask, the relevant traders will not trade. In this case, the mechanism generates money, and so requires an individual to absorb this, and this failure to be *budget-balanced*[33] is the price one has to pay for the strategy-proofness.

[32] As McAfee [43] points out, the advantage of the BBDA is that since sellers always tell the truth, it is easier to analyse this kind of auction than, for example, the CH or CDA.

[33] A mechanism is budget-balanced if the only flow of money is from one trader to another — a mechanism that is not budget-balanced either generates surplus for individuals other than the traders, or, worse, requires the addition of money.

4.2 Trader behaviour

The main point of interest with respect to traders is how they decide what offers to make, that is "price formation". Some of these models that have been developed to explain price formation are purely theoretical, some are validated against observed behavior in auction experiments, and yet others are used as the basis for trading agents. The results we consider here are the more purely theoretical, and those used as the basis for trading agents are considered in more detail in Section 5.

These results on price formation are important because, while we have, as above, results for what happens at equilibrium, as Satterthwaite and Williams point out [59]:

> no well-developed theory exists that explains how traders starting *de novo* jointly learn a set of equilibrium strategies (especially of more than one equilibrium exists). In other words, the theory does not include a process that results in . . . equilibrium.

Mechanisms for price formation are an attempt to fill this gap.

The earliest work on price formation by traders seems to have been that of [18], credited by [3] but about which we have been able to find no further information. This was followed by the model proposed by Easley and Ledyard [11][34] which establishes a series of plausible heuristics. Under this model, traders start by establishing a reservation price for their goods (which can differ from their private value), and then adjust this reservation price as the auction proceeds. The model is constructed in terms of assumptions about the reservation prices rather than an algorithm for computing them, but yield (as we shall see shortly) predictions about offers that can be tested.

According to [3], Friedman's "Bayesian game against nature" (BGAN) model [13] has the following characteristics. First, all traders formulate a "reservation price" for the good in question, a price that is an estimate of a reasonable price at which to trade, and will not trade unless it is at a profit with respect to the reservation price. Second, a buyer will make a bid if the highest bid is lower than its reservation price, increasing that bid slightly ([3] say "a buyer seizes the market bid", which we take ot mean slightly betters it), and if the lowest ask allows a profitable trade, a BGAN trader will accept it. Sellers have symmetrical behavior. Third, traders update their beliefs about possible offers by other traders in a Bayesian way, ignoring any effect that their own offers will have — it is this latter aspect that makes the trader behavior a "game against nature" — using this to set their reservation prices.

Another early model of price formation is that suggested by Wilson [76], and called the "waiting game/Dutch auction" (WGDA) in [3], from which this description is drawn. As the name suggests, this model has buyers and sellers

[34] This paper seems to have languished for some time before being published in [15], which explains how Friedman and Cason [3] can suggest the above chronology while [11] makes reference to [13] as an extant model.

initially doing nothing, waiting to see if another trader makes an offer[35]. When an offer is made, ending the waiting game, it may of course be accepted immediately. If it is not, the trader that made the offer is assumed to operate as the auctioneer in a Dutch auction, raising its bid, or lowering its ask, until its offer is finally accepted. During the Dutch auction phase, other traders will not make offers, and the remaining traders resume the waiting game at the end of the auction. A key aspect of the model is the notion that each trader has a level of impatience, which drives them to end the waiting game by making their initial offer, and this is also used to set the transaction price — the transaction price for a given trade splits the profit from that transaction between buyer and seller according to the ratio of those buyers and sellers that have not yet traded.

We can also consider Gode and Sunder's zero-intelligence traders, discussed above, as a theory of price formation, though it is something of a stretch to consider the random bidding of such agents as a true theory of price formation, and it is in the spirit of such theories that Gode and Sunder analyse the theoretical performance of these traders [21] in different kids of auction. The initial case that Gode and Sunder analyze is a *synchronised* double auction — intended as an abstract description of the market from [54] — in which all traders make offers and there is a trade if the highest bid and lowest ask cross. If there is no trade, new offers are solicited, subject to the rule that the highest bid cannot be reduced and the lowest ask cannot be increased[36] and a trade cancels all unaccepted offers. In this market ZI-C traders will have a minimum expected efficiency of 80.84%, and the average expected efficiency is computed to be around 96%.

This analysis is followed in [21] with an analysis of an unsynchronised market. This market is created by randomly soliciting offers from traders, and allowing the market to clear as soon as a bid and ask cross, and more offers are solicited until a trade takes place. As before, trades cancel outstanding bids and asks. Note that offers are solicited without replacement, so every trader gets to make one offer before any other trader can make two offers. In this market, the minimum expected efficiency falls to 75%, suggesting that, in general, continuous auctions may see lower efficiencies than call markets.

The three theories for price formation that we have just covered, those of Friedman [13], Wilson [76], and Gode and Sunder [20] are evaluated in [3]. This evaluation identifes testable hypotheses made by the theories — for example hypotheses about the order in which trades should occur and the correlation between prices — and identifies which of these are consistent with data from laboratory experiments with human subjects. In other words, the paper tests the extent to which these theories provide an adequate explanation of the way

[35] One of the predictions mde by the model is the waiting game is such that early transactions will be between buyers with high provate values and sellers with low private values, and later transactions will be between buyers with lower private values and sellers with higher private values.

[36] The mechanism for this is not clear in [21], and at least two approaches can be imagined. In one the traders with those limiting offers are not allowed to replace them with worse offers, while in the other, traders are repeatedly solicited for offers until the previous limiting bids are exceeded.

that human traders behave. The results are not very clearcut.

Thus the data suggests that all three theories correctly predict that efficiency will be close to 100%, and that all three correctly predict that buyers with high private values and sellers with low private values will tend to trade before other traders. However, Friedman's BGAN model, which suggests that traders look to improve on offers made by others, fits the data more closely than the others, while (rather surprisingly perhaps) the model that best fits the observed sequence of transaction prices is the ZI-C model.

Another theory of price formation is embodied in an algorithm proposed by Gjerstad and Dickhaut [19]. This approach, called GD in the literature, picks offers by computing the expected value of a range of offers, and then choosing the offer with the highest expected profit. The approach looks at all the offers that are competitive, that is all offers between the lowest unaccepted ask and the highest unaccepted bid, and computes an expected profit for each. The value part is easy to establish — it is just the value of the offer minus the private value — but the probability is more complex. GD computes the probability by looking at previous offers and considering the distribution of those offers as an estimate of the probability that a future offer will be accepted at the same price. In particular, GD fits a cubic to the offers in the space between the lowest unaccepted ask and the highest unaccepted bid and then, and uses the probabilities that this cubic generates to establish the expected profit of all the possible offers in that space. While somewhat complex to compute (especially the part that fits a cubic to the offers), GD (as we shall see below) comfortably outperforms other tarding mechanisms.

A somewhat similar approach to that of GD is that proposed by Roth and Erev [12, 53]. Here the agents also construct a probability distribution over the set of possible offers, but do this by reinforcement learning — basically trying the offers and seeing if they are accepted. Learning like this takes longer ([46] mentions experiments over 1000 trading periods) than for GD, but closely mimics human learning in markets [12, 53]. A modified version of the Roth-Erev algorithm (MRE) is used by Nicolaisen *et al.* [46], who also describe in detail some of the behaviours of the original algorithm that led to their modifications. The MRE approach is designed to work in call markets, and does not seem to perform anywhere near as well in continuous double auctions.

There are many other models of price formation and we do not have room to include them all here. Those that seem particularly worth mentioning include the various entries in the Santa Fe tournament [54, 55] each of which embody such a model. Some of the more interesting entries are as follows:

- Kaplan is the strategy that won the tournament. It avoids making offers, perferring to wait until an offer is made by another trader that Kaplan can accept at a good profit[37].

[37] If one considers the making of bids and asks that do not match, but which are adjusted towards each other over some period of time, as the negotiation of a deal, then Kaplan can be described, as it is in [55] as "wait in the background and let others do the negotiating, but when bid and ask get sufficiently close, jump in and 'steal the deal"'.

- Ringuette is a strategy that is very similar to Kaplan, though developed independently, and came in second in the tournament.
- GAMER is a simple strategy that always makes offers that yeld a 10% profit.
- Truthteller simply makes offers at its private value.
- Max operates in a very similar way to GD, using past offers to compute the probability of an offer being accepted and then computing the expected value of possible offers[38].

These, and other, strategies are detailed in [55][39]. Other notable approaches to price formation include Cason and Friedman's "Nash robots" [4], which use the computed Bayesian Nash equilibrium of a market (in the case of [4], a call market in which offers are sealed) to decide what offers to make, and Madhavan's models of price formation in quote-driven and order-driven securities markets [38], in the opening of a market by a specialist [40], and in more general specialist markets [36].

We close this section with a little theory of our own. In [21], Gode and Sunder do not discuss the performance of ZI-U agents, but it is straightforward to obtain such results, and they make a nice contrast with the results for ZI-C agents. First, it is easy to establish the profit made by a market in which just ZI-U traders are operating:

Lemma 1 *Consider a market of n buyers and m sellers, all of which bid using the* ZI-U *strategy, where b_i and s_j are the private values of the buyer i and seller j respectively, and the traders are ordered so that for all i and j, $b_i \geq b_{i+1}$ and $s_j \leq s_{j+1}$.*

If $n > m$, and trading continues until all sellers have traded, then the surplus S_a generated by the market is such that:

$$\sum_{i=(n-m)+1}^{i=n} b_i - \sum_{j=1}^{j=m} s_j \leq S_a \leq \sum_{i=1}^{i=m} b_i - \sum_{j=1}^{j=m} s_j$$

If $n \leq m$, and trading continues until all buyers have traded, then the surplus S_a generated is such that:

$$\sum_{i=1}^{i=n} b_i - \sum_{j=1}^{j=(n-m)+1} s_j \leq S_a \leq \sum_{i=1}^{i=n} b_i - \sum_{j=1}^{j=n} s_j$$

Proof: *The proof is straightforward. Consider the case where $n > m$. With no budget constraint on the traders, they bid randomly. Given long enough,*

[38] Due to a what [55] identifies as a bug in the code, Max did poorly in the tournament.

[39] Traders based on BGAN [13] and Easley and Ledyard's heuristic rules [11] were also entered.

eventually every seller will make an offer that is less than or equal to that made by some buyer, and the profit from the trade between the i*th buyer and the* j*th seller will be:*

$$b_i - s_j$$

which, of course, depends only on the private values of the traders, not the offers they made. The total surplus is just this value summed across all transactions. The precise value will depend upon which of the excess buyers traded. The maximum value of the surplus will be when the buyers with the highest private values trade (the buyers with the lowest m *values of* i*), and the lowest value of the surplus will be when the buyers with the lowest private values trade (the buyers with the highest* m *values of* i*). These bounds on the surplus give the first part of the result, and the same line of reasoning when* $n \leq m$ *gives the second part of the result.* □

An immediate corollary of this is that if the numbers of buyers and sellers is the same, then if the market continues until trading ceases, there is no variance in surplus, and hence in efficiency, over many iterations of the same market (as we will see in Section 5). Note also that this result means that the efficiency of an ZI-U market can be negative.

What efficiency this surplus will translate into obviously depends on the relationship between the private values of buyers and sellers. Making some assumptions about the range of private values — an assumption that we can consider defines a sensible market, one in which the supply and demand curves cross — we can show that the efficient surplus of the market, the surplus if trading is 100% efficient, is:

Lemma 2 *Consider a market of* n *buyers and* m *sellers, all of which bid using the* ZI-U *strategy, where* b_i *and* s_j *are the private values of the buyer* i *and seller* j *respectively, and the traders are ordered so that for all* i *and* j, $b_i \geq b_{i+1}$ *and* $s_j \leq s_{j+1}$*. Let* i^* *denote the largest value of* i *such that* $b_i \leq s_j$ *for some* j*, and* j^**, denote the smallest value of* j *such that* $s_j \geq b_i$ *for some* i*.*

If $i^* \leq j^*$*, the efficient surplus* S_t *is:*

$$S_t = \sum_{i=1}^{i=i^*} b_i - \sum_{j=1}^{j=i^*} s_j$$

and if $i^* > j^*$*, the efficient surplus* S_t *is:*

$$S_t = \sum_{i=1}^{i=j^*} b_i - \sum_{j=1}^{j=j^*} s_j$$

Proof: *The proof is very easy. The supply and demand curves cross at the* i^**th buyer and the* j^**th seller. If there are less buyers than sellers on the left hand side of this crossing point, then all the buyers trade at equilibrium, and* *so do the first i^* sellers. (We are assuming that the supply and demand curves cross). If there are more buyers than sellers to the left, then the first j^* buyers and all the sellers will trade at equilibrium.* □

and from Lemmas 1 and Lemma 2, we can obtain the overall efficiency of ZI-U traders:

Proposition 1 *Consider a market of n buyers and m sellers, all of which bid using the ZI-U strategy, where b_i and s_j are the private values of the buyer i and seller j respectively, and the traders are ordered so that for all i and j, $b_i \geq b_{i+1}$ and $s_j \leq s_{j+1}$. Let i^* denote the largest value of i such that $b_i \leq s_j$ for some j, and j^*, denote the smallest value of j such that $s_j \geq b_i$ for some i.*

If $n < m$ and $i^* \leq j^*$ the allocative efficiency α_{ziu} of the market will fall in the range:

$$\frac{1}{\left(\sum_{i=1}^{i=i^*} b_i - \sum_{j=1}^{j=i^*} s_j\right)} \left[\left(\sum_{i=(n-m)+1}^{i=n} b_i - \sum_{j=1}^{j=m} s_j \right), \left(\sum_{i=1}^{i=m} b_i - \sum_{j=1}^{j=m} s_j \right) \right]$$

Proof: *Direct from Lemmas 1 and 2.* □

Results for the symmetrical cases can easily be obtained.

Madhavan [38] gives a result that shows that a large enough call market can give more efficient prices than a continuous double auction, a result that agrees with Gode and Sunder's result for ZI-C agents, where lower minimum efficiency is expected in a continuous market. Furthermore, it is easy to obtain the following result, which supports both the above, independent of the strategy agents use to make offers or the size of the market:

Proposition 2 *For a given set of offers from agents that do not make loss-making offers, if the offers are made in random order, the expected surplus generated by a CH is greater than the expected surplus generated by a CDA.*

Proof: *Consider the bids and asks from Figure 1 being made in a CH. A CH, which has the luxury of being able to draw up supply and demand schedules, will match the highest bid with the lowest ask, the second highest bid with the second lowest ask, and so on. Let us imagine that the surplus so generated is S — this is clearly the greatest surplus that can be generated from the given offers. It will, given the fact that agents do not make loss-making offers, never be an over-estimate of the total possible surplus.*

Now, consider that the same offers being made in a CDA. If the offers arrive such that the highest bid is made first and is directly followed by the lowest ask, and then the second-highest bid arrives, directly followed by the second-lowest ask, and so on, then the CDA will generate surplus S, and there are other orders of arrival that will also generate S (the two highest bids arrive, followed by the two lowest asks followed by the same sequence as before, for example). However, *it is also clear that there are sequences for which the surplus generated will be less than* S*, and so, over all possible sequences of offers, the average surplus will be less than* S*. Thus the expected surplus is less than* S*, and the result holds.* □

Note the restriction that traders have to bid rationally — this ensures that S is always an underestimate of the true surplus, preventing the CDA generating higher surplus than the CH by not matching loss-making offers. In addition, note that the condition that offers be made in random order is stronger than necessary. The result will hold so long as offers are not only made in an order that always generates S. Further, note that the larger the market, the more likely it is that the continuous market will generate a lower surplus (since there are more sequences of offers that will not be exactly the right sequence to generate the same surplus S as the call market.

5 Experimental work on double auctions

Having categorised some of the auctions that have been studied in the literature, we discuss the experimental work that has been performed using some of those auctions in more detail. This leads on to some analysis of the way in which one can quantify auction performance.

5.1 A brief survey of experimental results

The work that started the experimental study of double auctions is that of Vernon Smith, and the first paper on this work is [61]. We have already described this work in considerable detail above, and it only remains to say, by way of comparison with the work described below, that it deals only with human traders (as one might imagine given when it was carried out), and examines how well markets work when they only have:

> . . . a practical number of marketers . . .

rather than the "indefinitely large" number previously assumed necessary[40]. The study concludes that even small numbers of traders allow competitive equilibrium to be achieved. This is demonstrated by the propensity of the transaction price to converge to the theoretical equilibrium price (the price at which the supply and demand curves, defined by trader private values, cross). These results held for a number of different supply and demand curves, suggesting a certain degree of robustness to the results.

Of course, this lone paper reflects just a very small fraction of Smith's work on experimental auctions (though, as the first, it may well be the most influential). Others include [35, 62, 63, 64, 65]. However, [61] is prototypical of much of the experimental work in economics, such was Smith's influence on the field,

[40] And showing that results like [75], discussed above, imply rather slower convergence than is actually the case.

Traders	Market 1	Market 2	Market 3	Market 4	Market 5
ZI-U	90.0	90.0	76.7	48.8	86.0
ZI-C	99.9	99.2	99.0	98.2	97.1
Human	99.7	99.1	100.0	99.1	90.2

Table 1: Efficiency results from [20].

and provides enough variation to keep us busy—if we can test artificial agents under all the conditions considered in [61] we will be doing very well indeed.

Other interesting experimental work with human subjects is that of Kagel and Vogt [32], who are concerned not with the true continuous double auction, like Smith, but with the buyer's bid double auction BBDA described above. Their aim was to test whether the market converges to 100% efficiency as it grows, as predicted by [59, **?**], and they did this by running examples of the BBDA with two buyers and two sellers and with eight buyers and eight sellers. The results of these experiments are compared with theoretical values (assuming the subjects bid in accordance with the Bayesian-Nash equilibrium) and the prediction made by ZI-C traders[41]. The human traders are somewhat less efficient than the theory predicts, but show a suitable increase in efficiency as the market gets larger.

The ZI-C traders do rather badly compared with their performance in CDA markets, attaining around 30% efficiency in the smaller market and up to 40% in the larger market. However, this seems likely to be because each trader only gets to make one offer in each trading period (the fact that they are sealed bids, of course makes no difference to ZI-C traders since they ignore offers made by other traders when choosing what to offer)—in CDA markets ZI-C traders typically make several offers before attaining a match.

[4]

One interpretation of Smith's work is that the kind of market provided by the double auction is strong enough that it manages to generate efficient outcomes even when assumptions like an "indefinitely large" number of traders are violated. Gode and Sunder [20] built on this work by investigating whether the market could cope with traders who made unintelligent offers rather than the rational traders assumed by theory. In particular, they looked at two varieties of what they termed *zero-intelligence* agents. Both of these bid randomly. ZI-U (unconstrained) agents pick a bid or ask from anywhere within a given price range (up to 200 monetary units in Gode and Sunder's experiments) and ZI-C (constrained) agents pick offers constrained by private value. Thus ZI-C buyers bid below their private value and ZI-C sellers ask above their private value, ensuring that they trade at a profit.

The results of this investigation were that while prices did not converge to equilibrium (which one would not expect given the traders are picking prices

[41] [32] calls these "ZI traders", but makes it clear that these are "budget-constrained zero-intelligence traders.

Traders	Market 1	Market 2	Market 3	Market 4	Market 5
ZI-U	225.48	253.12	90.54	363.80	156.28
ZI-C	28.53	49.81	15.90	60.47	19.07
Human	18.67	28.74	8.23	15.37	30.69

Table 2: Profit dispersion from [20].

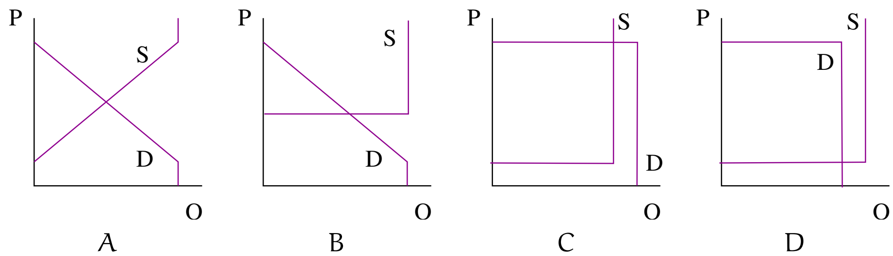

Figure 3: The four types of market defined by Cliff.

randomly) they did trade close to the theoretical equilibrium price, and the markets with ZI-C agents had high efficiency. This is a measure of how effective the market is at generating profits for the traders. For each trade it is possible to determine the profit for each trader (the difference between the trade-price and their private value) and so to establish the total actual profit for the auction. This can then be use to compute the fraction of the total possible profit that the auction could have extracted—if the buyer with the highest private value trades with the seller with the lowest private value, and the buyer with the second highest private value trades with the seller with the second lowest private value, and so on until the private values no longer cross then the maximum possible profit is extracted. The human markets (and Gode and Sunder ran some experiments with human markets with the same private value distributions for comparison) traded at around 100% efficiency and ZI-U traders at between 48.8% and 90%, while ZI-C traders achieved between 96 and 99.9% (averaging 98.7% across the multiple periods and the different private value distributions)[42].

Gode and Sunder also computed the *profit dispersion* for their experiments (as defined above). The profit dispersion figures are given in Table 2. These figures show that while ZI-U agents have a high profit dispersion, ZI-C agents have a profit dispersion that is much closer to that of human traders.

Taken together, these results would seem to suggest that traders don't have to have any real intelligence. at all for the market work well. All that traders need to do is to avoid trading at a loss[43], a restriction Gode and Sunder refer to

[42] These results are replicated in Table 1

[43] As we showed above, ZI-U traders can achieve large negative efficiency if the trades are such that both buyer and seller trade at a loss, though the structure of the markets used by Gode and Sunder largely rules this possibility out.

d	Market A	Market B	Market C	Market D
1	201 (12)	233 (09)	138 (14)	249 (14)
2	210 (14)	234 (10)	137 (17)	246 (13)
4	199 (14)	236 (10)	135 (17)	253 (12)
6	199 (13)	232 (10)	136 (16)	251 (12)
10	201 (15)	235 (10)	137 (15)	248 (13)
30			135 (16)	250 (12)

Table 3: Mean transaction prices from [7] for trading period d using ZI-C. Standard deviations are given in parentheses.

d	Market A	Market B	Market C	Market D
1	35 (09)	27 (08)	62 (17)	48 (11)
2	33 (09)	28 (07)	61 (13)	48 (10)
4	32 (09)	27 (08)	64 (15)	52 (10)
6	33 (10)	26 (07)	60 (15)	51 (08)
10	33 (09)	27 (07)	58 (12)	49 (13)

Table 4: Profit dispersion from [7] for trading period d using ZI-C. Standard deviations are given in parentheses.

as "market discipline", the idea being that traders who don't do this will rapidly go bankrupt and be forced out of the market. The double auction mechanism itself seems to be all that is required for efficient trading.

However, this result has subsequently been challenged. Cliff [6, 7, 8, 9] suggests that ZI-C agents trading close to the theoretical equilibrium has more to do with the sets of private values chosen than the power of the market. For a classic upward-sloping supply curve (more units are sold as the price increases) and a downward-sloping demand curve (less units are bought as the price decreases) of Figure 1, Cliff argues that the mean of the probability distribution from which the ZI-C agents pick prices is around the theoretical equilibrium price. Thus it is no surprise that for such markets, which are exactly the kind of markets studied in [20] (and most of the markets studied in [61]), ZI-C agents trade relatively close to the equilibrium price. This effect is exacebated by the fact that the markets used in [20] are such that there are very few *extra-marginal* traders—that is traders whose private values are such that they would not trade at equilibrium. In a "blind" market like that populated by ZI-C traders, every extra-marginal trader is a chance to lower efficiency and create trades that diverge from the eqiulibrium price.

Cliff then tested the hypothesis that it is the private values (which determine the shape of the supply and demand curves) that have the biggest effect on the apparent equilibrium price by running simulations of auctions with different supply and demand curves. In addition to the kind of markets studied by Gode and Sunder, Cliff used markets with:

```
if the last bid or ask was accepted at price P
then
    if V_B ≥ P then increase profit margin
    if the last offer was an ask
    then
         if V_B ≤ P then decrease profit margin
else
    if the last offer was a bid
    then
         if V ≤ P then decrease profit margin
```

Figure 4: A part of the ZIP algorithm.

- Downward sloping demand curve and flat supply curve.
- Flat supply and demand curve with excess demand (fewer sellers than buyers, and all sellers have a private value below that of the buyers).
- Flat supply and demand curves with excess supply (fewer buyers than sellers, and all sellers have a private value below that of the buyers).

Giving a set of markets with the shapes plotted in Figure 3 (the precise private values used in Cliff's experiments are given in [6, 9]).

The results confirm Cliff's hypothesis—ZI-C agents trade a significant distance away from the theoretical equilibrium in these new kinds of market. (In contrast, as shown by Smith [61, 62] — and in [63] as cited by [48] — markets with human traders show no such variation when demand and supply curves vary.) Note that Cliff reports no results about the efficiency of the ZI-C agents, and this need not be low just because trading takes place away from the theoretical equilibrium. To establish the allocative efficiency, we sum all the profit made by all the agents, and divide that by the total possible profit. Thus a buyer B and a seller S who trade at a price P when they have private values $\mathsf{V_B}$ and $\mathsf{V_S}$, contribute:

$$(\mathsf{V_B} - \mathsf{P}) + (\mathsf{P} - \mathsf{V_S})$$

to the numerator of this computation, and this, of course, reduces to:

$$\mathsf{V_B} - \mathsf{V_S}$$

so the efficiency depends only on the private values of the traders, not on the price at which they trade (providing the latter is between their private values). In other words, high allocative efficiency depends only upon matching buyers that have high private values with sellers that have low private values.

Instead of measuring efficiency, Cliff measured the average transaction price (for comparison with the transaction price in theoretical equilibrium, which is 200 for the markets he studied) and the same profit dispersion measure as in

d	Market A	Market B	Market C	Market D
1	185 (19)	225 (07)	126 (20)	240 (08)
2	191 (11)	208 (05)	129 (17)	237 (09)
4	196 (06)	201 (01)	138 (19)	228 (10)
6	199 (04)	201 (01)	143 (20)	222 (10)
10	200 (02)	201 (01)	152 (21)	213 (11)
30			186 (20)	203 (07)

Table 5: Mean transaction prices for trading period d using ZIP, from [7]. Standard deviations are given in parantheses.

d	Market A	Market B	Market C	Market D
1	27 (09)	20 (06)	64 (19)	38 (07)
2	13 (09)	06 (04)	58 (18)	34 (07)
4	06 (07)	01 (01)	51 (19)	27 (07)
6	04 (05)	01 (01)	47 (19)	20 (08)
10	03 (06)	01 (01)	40 (20)	10 (09)

Table 6: Profit dispersion for trading period d using ZIP, from [7]. Standard deviations are given in parantheses.

Gode and Sunder. These values are given in Tables 3 and 4. Clearly ZI-C does not converge well to the transaction price in markets B, C and D, and the profit dispersion is similar to that reported by Gode and Sunder (who give no standard deviations).

In addition to his critique of the Gode and Sunder's work, Cliff [6, 9, 7] provides an alternative approach, the *zero-intelligence-plus* (ZIP) trader. This trader listens to the bids and asks being made, and depending on how these relate to its private value, raises or lowers its bid (or ask). For example, a buyer with private value V_B operates as in Figure 4 where every buyer is bidding above its private value by its profit margin, and every seller is bidding below its private value by its profit margin. The increases and decreases are carried out using a variation of the Widrow Hoff rule [72]. For full details see [6, 9]. The results obtained for ZIP agents are given in Tables 5 and 6, where d is the number of the trading period a given set of values is from. It is clear from these figures that ZIP agents outperform ZI-C agents both in terms of how closely they approach the transaction price (especially when given enough trading periods to learn over) and in the levels of profit dispersion.

Preist and van Tol, working at HP labs, where Cliff carried out the work reported above, extended his work in [49]. They started with the observation that one of the assumptions underlying Cliff's work, that all outstanding bids and asks are cleared after every trade, is somewhat unrealistic for many markets. As a result, they investigated persistent shout double auctions, and adapted the ZIP algorithm to deal with this case. Their new approach, which we will refer to

d	Human	GD	ZI-C
1–2	0.907	0.9982	0.968
1–10	0.959	0.9991	0.968
9–10	0.970	0.9992	0.967

Table 7: Efficiency for trading period d, from [19].

as PVT, was tested against ZIP traders in one of the markets that Cliff introduced in [6] (the symmetric market Cliff calls "A"). Given that Cliff's interest was with the degree to which traders approach the theoretical equilibrium price, it is natural that [49] should look at this also, and they use the α measure to show that PVT agents provide a more stable solution (meaning that they have a lower α measure, a mean value of around 0.01 compared to 0.04 for ZIP) and they further claim that PVT traders approach equilibrium faster.

Both the PVT traders and the ZIP traders on which they are based use a relatively simple algorithm which tries to keep track of the best profit margin for the trader, where this notion of "best" tries to balance making a big profit with ensuring that an offer is accepted. These approaches make an interesting contrast with the GD and MRE approaches discussed above. Both the latter try to estimate the expected value of a given offer, and so explicitly deal with both the size of the profit margin and an explicit measure of how likely a given offer is to be accepted.

Gjerstad and Dickhaut, whose automated GD trader was described above, tested their approach on seven different markets[44], comparing the performance of GD with that of ZI-C and human traders (the figures for human traders being taken from [35]), and measuring both efficiency and the mean absolute deviation of the transaction price from the theoretical equilibrium price. Values for both GD and ZI-C were averaged over 100 simulations (no standard deviation data are given in the paper), and each simulation was run over 10 trading periods. Because one of the features that Gjerstad and Dickhaut were interested in was the ability of their model to respond to changes in the market, private values were changed after five periods — every agent's private value was increased by $0.50 (initial values being between between roughly $1.40 and $3.50). Results from these experiments are given in Tables 7 and 8, and show that GD outperforms both ZI-C and human traders in all cases. Furthermore, as in Gode and Sunder's original experiments, ZI-C has a slight edge over human traders in efficiency terms in early periods, while trading a little further from equilibrium.

All the work considered so far looked at auctions in which only one kind of agent took part (counting human traders as one kind of agent, as studies typically do, despite the fact that they doubtless use a range of strategies). It

[44]Markets named 3pda01, 2pda17, 2pda20, 2pda21, pda24, 2pda47, and 2pda53, which Gjerstad and Dickhaut state were originally investigated by Ketcham *et al.* [35], but which we cannot find described in [35]. Though one can assume they are versions of the markets xxxxx, this is unfortuately not enough information to allow us to recreate the markets.

d	Human	GD	ZI-C
1–2	0.101	0.077	0.276
1–10	0.050	0.045	0.237
9–10	0.022	0.040	0.209

Table 8: Mean absolute deviation from equilibrium price for trading period **d**, from [19].

is also interesting, and arguably more realistic, to investigate auctions in which traders use a variety of bidding strategies.

An interesting example of this line of work is [10] (also discussed in [34]), which pits human traders against trading agents. Prior to this work, it had been established that human traders quickly learnt to trade near equilibrium, and could outperform ZI-U and ZI-C traders. However, nobody had compared human performance with that of the "more than zero" intelligence traders that had been developed. [10] pitted 6 humans against 6 agents — both types of trading agent were both buyer and seller in all these experiments — where the agents used modified versions of GD (MGD) and ZIP (MZIP), the modification being necessary to deal with an order book that was not cleared at every trade[45]. MGD is described in some detail, MZIP is said to be similar to PVT. Results, albeit only on six experiments, show that agents tend to be more efficient traders (agents managed to trade at over 100% efficiency, suggesting that they steal profit from the human traders, while the humans traded between 65% and 97% efficiency).

The idea of testing the performance of a mixture of human and automated traders is also explored in [23], though in a rather different market. The market is modelled on the US futures market, with one seller and five buyers, and the automated traders are *arbitrageurs* that exploit the difference between market prices and those offered by an external bank. The main aspect tested in the experiment described in [23] was that of the effect of information on traders. Results show that providing information on the automated traders to human traders helps the convergence to equilibrium, and in cases when there is no such information, markets with automated traders operate at lower efficiency than those without such agents.

Comparing a range of automated trading strategies was the aim of the Santa Fe Double Auction Tournament [54, 55]. In order to encourage a variety of entries, the organisers offered cash prizes and attracted 30 submissions. These entries were then played against one another, with one copies of each entry

[45] [34] erroneously suggests that ZIP needed to be modified because it was initially intended for a call market, but close reading of [6] shows this to be false, and [10] confirms that the main modification is to deal with outstanding orders. [34] also includes the intriiguing statement that "We have already begin to develop a stringer bidding algorithm, GDX, that resoundingly beats GD and ZIP. Informal tests indicate that GDX fares even better against humans than GD", but we are unware of any other reference to GDX in the literature.

pitted against copies of other entries in a series of different types of market[46]. In this respect the tournament was the complete opposite of the homogeneous set of traders in both Cliff's and Gode and Sunder's experiments. The aim was also different — here the focus was on finding the strategy which gains the greatest profit over the course of many runs of the market (and thus from different private values).

The two top strategies in the tournament were both *snipers*[47], that is strategies which wait until a deal seems settled and then run in and steal it. The top scoring program, named *Kaplan* in honour of its creator, is also very simple. It does not bid until the gap between the latest bid and ask is less than 10%, and then bids equal to that last ask as long as it is profitable to do so. Such a simple strategy is surprisingly effective, and, as discussed in [54], even manages to make good profits when bidding against large number of copies of itself. This latter fact was established through an evolutionary tournament like that of Axelrod [1] in which successful strategies were allowed to "breed" and unsuccessful ones were eliminated, though this tournament also revealed that Kaplan traders are not stable — eventually the population of Kaplan traders goes into decline.

The program that came second, Ringuette, was only just beaten by Kaplan, ending the tournament with a payff of \$394 rather than \$408[48]. According to [55], the payoff distributions for the two strategies are "almost identical" and the main difference between the two is that Ringuette sometimes trades at a loss, and sometimes fails to make profitable trades. In particular, Ringuette is reported as trading 99% of the items it should be able to trade profitably (which we assume means given the theoretical equilibrium of the various scenarios it faces) while Kaplan trades 111%. [55] attributes this to Ringuette using a "more aggresive" rule to decide when to make an offer, and it randomly overbids an ask it aims to match, whereas Kaplan just matches it.

One reasonable conclusion that might be drawn from the Santa Fe tournament is that sniping strategies do well, but such a conclusion should be viewed with caution. Snipers require other strategies to snipe off — as we have already pointed out, Kaplan cannot continue to earn good profits if all other traders use Kaplan (even though the strategy includes a mechanism that is intended to ensure it trades if there are no traders to snipe at, a strategy that involves starting bidding as time runs out in the auction). More correct would be to say that sniping strategies can do well, provided that they operate in a suitable mix of strategies[49]. Indeed, one of the main conclusions one can draw from all experimental work on trading strategies is that the performance of any strategy

[46] In a complicated arrangement that tried to ensure that different entries were treated as identically as possible.

[47] A term not in use at the time the experiments were carried out, but familiar to anyone who uses eBay.

[48] Though this is a small margin, [55] points out that the difference in mean efficiency between the two strategies is significant.

[49] As discussed above, [54] suggests that the Kaplan strategy does not give terribly good results in markets in which it is the only strategy to be used, and this is backed up by results from our own experiments. These results contrast with those reported in [66, 67] and discussed below.

tends to be conditional on the other strategies it operates against.

A further conclusion is that strategies that perform well do not need to be complex. While Kaplan is not the simplest strategy that took part in the tournament — GAMER, for example, with its constant profit margin, is simpler and and performs poorly — it is much simpler than, for example, a neural network trader that was entered, which included 20 hidden units and 40 inputs from various pieces of public information (and clearly does better).

Work in a similar vein to the Sante Fe tournament has recently been carried out by Tesauro, Das and colleagues. In the search for traders which are capable of making high profits, Tesauro and Das [66, 67] exhaustively tested five strategies —ZI-C, Kaplan, ZIP, GD, and MGD, a modified version of GD. The results show that homogeneous sets of traders of all strategies gain high efficiency (the highest is 0.997 for ZIP and MGD and the lowest is 0.983 for ZI-C), that both ZIP, MGD converge quickly to an equilibrium (results for GD are not given), and that these latter two also come very close to the theoretical equilibrium price on average. None of these results should be surprising given the right type of market[50].

These results, of course, only measure the working of the market as a whole. To find the effectiveness of strategies for individual traders in the same vein as the Santa Fe experiments, Tesauro and Das carried out further experiments. First they tested whether traders in an otherwise homogeneous population gained or lossed by changing strategy. This showed that it was never profitable to change to using ZI-C from any of the other strategies, and that it is always better to switch from ZI-C. Similarly traders can always benefit by switching to Kaplan, but this benefit is small against GD and MGD traders. A single ZIP trader always does well, except against MGD, and MGD always does well, except against ZIP (and no results are given against GD).

With more balanced groups of types of agent, a number of direct comparisons were made. Over 100 runs, the number of "wins" (where a given group of traders gained more profit than the other group) was counted. This showed that ZI-C could defeat Kaplan, GD could defeat ZI-C and Kaplan, ZIP could defeat GD, ZI-C and Kaplan, and MGD could defeat all the others (and by a significant margin). This leds Tesauro and Das to the conclusion that MGD dominates the other trading strategies, and, to some degree answers the question of what would have happened if GD had been entered into the Santa Fe tournament. It also reinforces the point, made above, that the performance of a trading strategy in an auction is dependent on the other stratgies that are used in the same auction.

A more sophisticated analysis, which extends the work of Tesauro and Das, is the heurstic strategy analysis of [69], work that is a natural extension of that in [66, 67][51]. The main idea of [69] is to apply evolutionary game theory [57] to the kinds of results that previous authors had obtained for trading agents. Overall the argument for the evolutionary approach is as follows. Consider a set of traders using different trading strategies, some of these will do better than

[50] Though see footnote 49.

[51] Which is not to say that it is unoriginal—the work was strikingly original, and only seems like an obvious extension in retrospect.

others, and which ones do well will depend upon the exact mix of traders—this is seem both in [66, 67] and the "breeding" work in [54]. Now, imagine a particular trader, using a specific strategy, after some number of auctions can identify how well traders using different strategies are performing in the current mix (taking profit per trader using a given strategy as a measure of the performance of that strategy). A rational trader will choose to adopt a different strategy if that strategy is performing better, and if there are several such strategies, will pick the one that is performing best. This will change the mix, and after some more auctions the trader may again change its strategy. If we establish the relative performance of different approaches for all the possible mixes (for some finite number of agents) we can determine whether there are stable equilibrium points that the set of traders will converge on, and what the mix is at such equilibria.

Of course, there is a huge number of possible strategies, and hence possible mixes, that we need to consider. The second key idea in [69] is that rather than choose an arbitrary strategy, sensible traders will stick with strategies like GD which are known to perform well. They choose to concentrate on Kaplan, ZIP and MGD, and looked at versions of the CDA with 6, 12, 14, 16, 18 and 20 agents. Results show that in the 6, 12 and 14 agent CDA, the Nash equilibrium is a mixed strategy in which MGD has zero probability (in other words MGD is outperformed by a mixture of ZIP and Kaplan), while in the 20 agent game there are three Nash mixed-strategy equilibria, two of which have MGD is the most likely strategy to be played. Furthermore, altering the payoffs slightly (to simulate an improved strategy) suggests that a modest improvement in the payoff to MGD would result in it becoming nearly dominant. Thus, it is clear that the precise results one obtains from experimental work on trading straegies seems very sensitive to the experimental conditions.

This last point is exacebated [70] by the fact that the computation of the payoff to an agent playing a given strategy is extremely computational intensive — it requires the running of many iterations of an auction, each with a considerable number of agents participating in it. To help reduce the computation, and to reduce the error in payoff calculation, error that will always be present to some extent given the natural statistical variation in the strategies that the agents employ, [70] suggest smart strategies to determine which auctions to run in order to compute the payoffs.

5.2 Learning

5.3 What existing experimental work does and does not tell us

6 Summary

This paper has attempted to collect and relate the bulk of the work on trading in double auction markets. The main aim of the paper was to start in some sense to unify the work that has been carried out in this area. That work is very diverse, and our aim is, by recording this diversity, to help build an understanding of not only the patchwork quilt of results but also the gaps in that quilt. This understanding is a necessary preliminary to our plan to provide the missing pieces of the quilt, running experiments that make the comparisons that are currently missing, and taking the measurements that allow such comparisons to be made. Some initial steps in this latter line of work were also described.

References

[1] R. Axelrod. *The evolution of cooperation.* Basic Books, New York, NY, 1984.

[2] E. G. Burrows and M. Wallace. *Gotham: A History of New York City to 1898.* Oxford University Press, New York, NY, 1999.

[3] T. N. Cason and D. Friedman. An empirical analysis of price formation in double auction markets. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, chapter 1, pages 253–283. Perseus Publishing, Cambridge, MA, 1993.

[4] T. N. Cason and D. Friedman. Price formation in single call markets. *Econometrica*, 65(2):311–345, March 1997.

[5] K. Chatterjee and W. Samuelson. Bargaining under incomplete information. *Operations Research*, 31(5):835–851, 1983.

[6] D. Cliff. Minimal-intelligence agents for bargaining behaviours in market-based environments. Technical Report HP-97-91, Hewlett-Packard Research Laboratories, Bristol, England, 1997.

[7] D. Cliff and J. Bruten. Less than human: Simple adaptive trading agents for CDA markets. Technical Report HP-97-155, Hewlett-Packard Research Laboratories, Bristol, England, 1997.

[8] D. Cliff and J. Bruten. More than zero intelligence needed for continuous double-auction trading. Technical Report HP-97-157, Hewlett-Packard Research Laboratories, Bristol, England, 1997.

[9] D. Cliff and J. Bruten. Zero is not enough: On the lower limit of agent intelligencefor continuous double auction markets. Technical Report HP-97-155, Hewlett-Packard Research Laboratories, Bristol, England, 1997.

[10] R. Das, J. E. Hanson, J. O. Kephart, and G. Tesauro. Agent-human interactions in the continuous double auction. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, Seattle, USA, 2001.

[11] D. Easley and J. O. Ledyard. Theories of price formation and exchange in oral double auctions. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, chapter 1, pages 63–97. Perseus Publishing, Cambridge, MA, 1993.

[12] I. Erev and A. E. Roth. Predicting how people play games with unique, mixed-strategy eqilibria. *American Economic Review*, 88:848–881, 1998.

[13] D. Friedman. A simple testable model of double auction markets. *Journal of Economic Behavior and Organization*, 15(1):47–70, 1991.

[14] D. Friedman. The double auction institution: A survey. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, chapter 1, pages 3–25. Perseus Publishing, Cambridge, MA, 1993.

[15] D. Friedman and J. Rust, editors. *The Double Auction Market: Institutions, Theories and Evidence*. Santa Fe Institute Studies in the Sciences of Complexity. Perseus Publishing, Cambridge, MA, 1993.

[16] D. Friedman and J. Rust. Preface. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, pages 199–219. Perseus Publishing, Cambridge, MA, 1993.

[17] D. Fudenberg, M. M. Mobius, and A. Szeidl. Existence of equilibrium in large double auctions. Working paper, Department of Economics, Harvard University, Littauer Center, Cambridge, MA 02138, April 2003.

[18] G. B. Garcia. On the dynamics of prices to equilibrium in a double auction market. Working paper, University of Chicago Business School, 1980.

[19] S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behaviour*, 22:1–29, 1998.

[20] D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial sustitute for individual rationality. *The Journal of Political Economy*, 101(1):119–137, February 1993.

[21] D. K. Gode and S. Sunder. Lower bounds for efficiency of surplus extraction in double auctions. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, chapter 7, pages 199–219. Perseus Publishing, Cambridge, MA, 1993.

[22] D. K. Gode and S. Sunder. Double auction dynamics: Structural effects of non-binding price contracts. Yale ICF Working Paper 00-17, Yale International Center for Finance, May 2003.

[23] J. Grossklags and C. Schmidt. Artificial software agents on thin double auction markets — a human trader experiment. Working paper, School of Information Management and Systems, University of California, Berkeley, 102 South Hall, Berkeley, CA 94720-4600, December 2002.

[24] L. Harris, G. Sofianos, and J. E. Shapiro. Program trading and intraday volatility. *The Review of Financial Studies*, 10(4):1035–1064, Winter 1997.

[25] J. Hasbrouck and G. Sofianos. The trades of market makers: An empirical analysis of NYSE specialists. *Journal of Finance*, 48(5):1565–1593, December 1993.

[26] J. Hasbrouck, G. Sofianos, and D. Sosebess. New York Stock Exchange systems and trading procedures. Working Paper 93-01, New York Stock Exchange, April 1993.

[27] B. Holmstrom and R. B. Myerson. Efficienct and durable decision rules with incomplete information. *Econometrica*, 51(6):1799–1820, 1983.

[28] L. Hurwicz. On informationally decentralized systems. In C. B. McGuire and R. Radner, editors, *Decision and Organisation: A Volume in Honour of Jacob Marchak*, pages 297–336. North-Holland, 1972.

[29] M. O. Jackson and J. M. Swinkels. Existence of equilibrium in single and double private value auctions. Mimeo, California Institute of Technology, 2001.

[30] M. O. Jackson and J. M. Swinkels. Existence of equilibrium in single and double private value auctions. *Econometrica*, 73(1):93–139, 2005.

[31] O. Kadan. Equilibrium in the two player, k-double auction with affiliated private values. Technical Report 12.2004, Nota di Lavoro, 2004.

[32] J. H. Kagel and W. Vogt. Buyer's bid double auctions: Preliminary experimental results. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, chapter 10, pages 285–305. Perseus Publishing, Cambridge, MA, 1993.

[33] D. B. Keim and A. Madhavan. The relation between stock market movements and NYSE seat prices. *Journal of Finance*, 55(6):2817–2840, December 2000.

[34] J. O. Kephart. Software agents and the route to the information economy. *Proceedings of the National Academy of Sciences*, 99(3):7207–7213, May 2002.

[35] J. Ketcham, V. L. Smith, and A. W. Williams. A comparison of posted-offer and double auction pricing institutions. *The Review of Economic Studies*, LI:595–614, 1984.

[36] J. C. Leach and A. N. Madhavan. Price experimentation and security market structure. *The Review of Financial Studies*, 6(2):375–404, 1993.

[37] G. Leffler and C. L. Farwell. *The Stock Market.* Ronald, New York, 1963.

[38] A. Madhavan. Trading mechanisms in securities markets. *The Journal of Finance*, 47(2):607–641, June 1992.

[39] A. Madhavan and M. Cheng. In search of liquidity: Block trades in the upstairs and downstairs markets. *The Review of Financial Studies*, 10(1):175–203, Spring 1997.

[40] 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, Autumn 2000.

[41] A. Madhavan, M. Richardson, and M. Roomans. Why do security prices change? A transaction-level analysis of NYSE stocks. *The Review of Financial Studies*, 10(4):1035–1064, Winter 1997.

[42] E. Maskin and J. Riley. Equilibrium in sealed high bid auctions. *The Review of Economic Studies*, 67(3):439–454, 2000.

[43] R. Preston McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56:434–450, 1992.

[44] K. A. McCabe, S. J. Rassenti, and V. L. Smith. Designing a uniform-price double auction: An experimental evaluation. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, pages 307–332. Perseus Publishing, Cambridge, MA, 1993.

[45] P. Milgrom and R. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50:1089–1122, 1982.

[46] J. Nicolaisen, V. Petrov, and L. Tesfatsion. Market power and efficiency in a computational electricity market with discriminatory double-auction pricing. *IEEE Transactions on Evolutionary Computation*, 5(5):504–523, 2001.

[47] J. Niu. Auction experiments with JASA. Technical report, Department of Computer Science, Graduate Center, City University of New York, 365 Fifth Avenue, New York, NY, May 2005.

[48] C. R. Plott. Industrial organization theory and experimental economics. *Journal of Economic Literature*, 20:1485–1527, 1982.

[49] C. Preist and M. van Tol. Adaptative agents in a persistent shout double auction. In *Proceedings of the 1st International Conference on the Internet, Computing and Economics*, pages 11–18. ACM Press, 1998.

[50] M. Raberto and S. Cincotti. Modeling and simulation of a double auction artificial financial market. *Physica A*, 355:34–45, 2005.

[51] P.J. Reny and M. Perry. Toward a strategic foundation for rational expectations equilibrium. Working paper, University of Chicago, 2003.

[52] D. J. Roberts and A. Postlewaite. The incentives for price-taking behavior in large exchange economies. *Econometrica*, 44(1):115–127, 1976.

[53] A. E. Roth and I. Erev. Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior*, 8:164–212, 1995.

[54] J. Rust, J. H. Miller, and R. Palmer. Behaviour of trading automata in a computerized double auction market. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, chapter 6, pages 155–199. Perseus Publishing, Cambridge, MA, 1993.

[55] J. Rust, J. H. Miller, and R. Palmer. Characterizing effective trading strategies. *Journal of Economic Dynamics and Control*, 18:61–96, 1994.

[56] A. Rustichini, M. A. Satterthwaite, and S. R. Williams. Convergence to efficiency in a simple market with incomplete information. *Econometrica*, 62(5):1041–1063, September 1994.

[57] L. Samuelson. *Evolutionary games and equilibrium selection*. MIT Press, Cambridge, MA, 1998.

[58] M. A. Satterthwaite and S. 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):477–498, 1989.

[59] M. A. Satterthwaite and S. R. Williams. The Bayesian theory of the k-double auction. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, chapter 4, pages 99–123. Perseus Publishing, Cambridge, MA, 1993.

[60] M. A. Satterthwaite and S. R. Williams. The optimality of a simple market mechanism. *Econometrica*, 70(5):1841–1863, 2002.

[61] V. L. Smith. An experimental study of competitive market behaviour. *The Journal of Political Economy*, 70(2):111–137, April 1962.

[62] V. L. Smith. Experimental auction markets and the Walrasian hypothesis. *The Journal of Political Economy*, 73(4):387–393, August 1965.

[63] V. L. Smith. Bidding and auctioning institutions: Experimental results. In Yakov Amihud, editor, *Bidding and auctioning for procurement and allocation*, pages 43–63. New York University Press, New York, 1976.

[64] V. L. Smith. Microeconomics as a experimental science. *The American Economic Review*, 72(5):923–955, December 1982.

[65] V. L. Smith and A. W. Williams. An experimental comparison of alternative rules for competitive market exchange. In *Auctions, Bidding and Contracting: Uses and Theory*. New York University Press, New York, 1982.

[66] G. Tesauro and R. Das. High-performance bidding agents for the continuous double auction. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*, pages 206–209, 2001.

[67] G. Tesauro and R. Das. High-performance bidding agents for the continuous double auction. In *Proceedings of the Workshop on Economic Agents, Models and Mechanisms*, 2001.

[68] W. Vickrey. Counterspeculation, auctions, and competitive sealed bids. *Journal of Finance*, 16(1):8–37, 1961.

[69] W. E. Walsh, R. Das, G. Tesauro, and J. O. Kephart. Analyzing complex strategic interactions in multi-agent systems. In P. Gymtrasiwicz and S. Parsons, editors, *Proceedings of the 4th Workshop on Game Theoretic and Decision Theoretic Agents*, 2001.

[70] W. E. Walsh, D. C. Parkes, and R. Das. Choosing samples to compute heuristic-strategy nash equilibrium. In P. Faratin, D. C. Parkes, J. A. Rodriguez-Aguilar, and W. E Walsh, editors, *Agent Mediated Electronic Commerce V*. Springer Verlag, 2004.

[71] M. P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–23, 1993.

[72] B. Widrow and M. Hoff. Adaptive switching circuits. In *IRE WESCON Convention Record*, volume 4, pages 96–104. 1960.

[73] S. R. Williams. The nature of equilibria in the buyer's bid double auction. Discussion Paper 793, Department of Economics, Northwestern University, 1988.

[74] S. R. Williams. Existence and convergence of equilibria in the buyer's bid double auction. *The Review of Economic Studies*, 58:351–374, 1991.

[75] R. Wilson. Incentive efficiency of double auctions. *Econometrica*, 53(5):1101–1116, 1985.

[76] R. Wilson. On equilibria of bid-ask markets. In G. Feiwel, editor, *Arrow and the Ascent of Modern Economic Theory*, pages 375–414. MacMillan, 1987.

[77] P. R. Wurman, W. E. Walsh, and M. P. Wellman. Flexible double auctions for electronic commerce: Theory and implementation. *Decision Support Systems*, 24:17–27, 1998.

[78] P. R. Wurman, M. P. Wellman, and W. E. Walsh. A parameterization of the auction design space. *Games and Economic Behavior*, 35(1/2):304–338, 2001.