

2009

TR-2009004: An Investigation Report on Auction Mechanism Design

Jinzhong Niu

Simon Parsons

Follow this and additional works at: http://academicworks.cuny.edu/gc_cs_tr

 Part of the [Computer Sciences Commons](#)

Recommended Citation

Niu, Jinzhong and Parsons, Simon, "TR-2009004: An Investigation Report on Auction Mechanism Design" (2009). *CUNY Academic Works*.

http://academicworks.cuny.edu/gc_cs_tr/325

This Technical Report is brought to you by CUNY Academic Works. It has been accepted for inclusion in Computer Science Technical Reports by an authorized administrator of CUNY Academic Works. For more information, please contact AcademicWorks@gc.cuny.edu.

An Investigation Report on

Auction Mechanism Design

April 7, 2009

Jinzhong Niu, Simon Parsons

Department of Computer Science

The Graduate Center, The City University of New York

Contents

1 Auctions	4
1.1 Auction types	4
1.2 Supply, demand and equilibrium	5
1.3 A typical time series of shouts	6
1.4 Performance metrics	8
1.4.1 Allocative efficiency	8
1.4.2 Convergence coefficient	9
2 Game theory	10
2.1 Games	10
2.2 Nash equilibrium	12
3 Auction theory	15
3.1 Revenue equivalence theorem	15
3.2 On double-sided auctions	16
4 Experimental approaches	19
4.1 Different data sources	19
4.2 Smith's experiments	20
5 Trading agents	24
5.1 Zero intelligence traders	24
5.2 Zero intelligence plus and beyond	25
5.3 Interaction of heterogenous trading strategies	29
5.4 Automating strategy acquisition	33
5.4.1 Evolutionary computation	33
5.4.2 Optimizing parameter combination in ZIP	34
5.4.3 Combining GA and heuristic strategy analysis	34
5.5 Trading Agent Competition	35

6	Experimental auction mechanism design	37
6.1	A parameterized space of auctions	37
6.2	Evolving market mechanisms	39
6.3	Evaluating market mechanisms	40
6.4	Adaptive auction mechanisms	41
6.5	Auction mechanism design competition	42
7	Summary	43

Abstract

Auctions are markets with strict regulations governing the information available to traders in the market and the possible actions they can take. Since well designed auctions achieve desirable economic outcomes, they have been widely used in solving real-world optimization problems, and in structuring stock or futures exchanges. Auctions also provide a very valuable testing-ground for economic theory, and they play an important role in computer-based control systems.

Auction mechanism design aims to manipulate the rules of an auction in order to achieve specific goals. Economists traditionally use mathematical methods, mainly game theory, to analyze auctions and design new auction forms. However, due to the high complexity of auctions, the mathematical models are typically simplified to obtain results, and this makes it difficult to apply results derived from such models to market environments in the real world. As a result, researchers are turning to empirical approaches.

This report aims to survey the theoretical and empirical approaches to designing auction mechanisms and trading strategies with more weights on empirical ones, and build the foundation for further research in the field.

1 Auctions

1.1 Auction types

A *market* is a set of arrangements by which buyers and sellers, collectively known as *traders*, are in contact to exchange goods or services. *Auctions*, a subclass of markets with strict regulations governing the information available to traders in the market and the possible actions they can take, have been widely used in solving real-world optimization problems, and in structuring stock or futures exchanges.

The most common kind of auction is the *English auction*, in which there is a single seller, and multiple buyers compete by making increasing bids for the commodity (good or service) being auctioned; the one who offers the highest price wins the right to purchase the commodity. Since only one type of trader—buyers—makes offers in an English auction, the auction belongs to the class of *single-sided auctions*. Another common single-sided auction is the *Dutch auction*, in which the auctioneer initially calls out a high price and then gradually lowers it until one bidder indicates they will accept that price.

Another class of single-sided auctions is the class of *sealed-bid auctions*, in which all buyers submit a single bid and do so simultaneously, i.e., without observing the bids of the others or if the others have bid. Two common sealed-bid auctions are the *first-price auction* and the *second-price auction* or *Vickrey auction* [47]. In both types of sealed-bid auctions, the highest bidder obtains the commodity. In the former, the highest bidder pays the price they bid, while in the latter, they pay the second highest price that was bid.

These four single-sided auctions—English, Dutch, first-price sealed-bid, and Vickrey—are commonly referred to as the *standard auctions* and were the basis of much early research on auctions.

In addition, there are *double-sided auctions* or *DAS*¹, in which both sellers and buyers make offers, or *shouts*. The two most common forms of DA are *clearing houses* or *CHs*² and *continuous double auctions* or *CDAs*. In a CH, an auctioneer first collects *bids*—shouts from buyers—and *asks*—shouts from sellers, and then clears the market at a price where the quantity of the commodity supplied equals the quantity demanded.

¹The terminology is not standardized, and sometimes these are called *bid-ask auctions*. Note that [14] used the term “double auction” to refer to what I call a continuous double auction in this report.

²These are sometimes called *call markets* or *static double auctions*.

This type of market clearing guarantees that if a given trader is involved in a transaction, all traders with more competitive offers are also involved.³ In a CDA, a trader can make a shout and accept an offer from someone at any time. This design makes a CDA able to process many transactions in a short time, but permits extra-marginal traders to make deals. Both kinds of DA are of practical importance, with, for example, CDA variants being widely used in real-world stock or trading markets including the New York Stock Exchange (NYSE) and the Chicago Mercantile Exchange (CME).

In some auctions, traders can place shouts on combinations of items, or “packages”, rather than just individual items. They are called *combinatorial auctions*. A common procedure in these markets is to auction the individual items and then at the end to accept bids for packages of items. Combinatorial auctions present a host of new challenges as compared to traditional auctions, including the so-called *winner determination problem*, which is how to efficiently determine the allocation once the bids have been submitted to the auctioneer.

Traders, in some cases, are allowed to both sell and buy during an auction. Such traders are called *two-way traders*, while those that only buy or only sell are called *one-way traders*.

This report will mainly discuss non-combinatorial DAs, especially CDAs, populated by one-way traders.

1.2 Supply, demand and equilibrium

A central concern in studies of auction mechanisms are the supply and demand schedules in a market. The quantity of a commodity that buyers are prepared to purchase at each possible price is referred to as the *demand*, and the quantity of a commodity that sellers are prepared to sell at each possible price is referred to as the *supply*. Thus if *price* is plotted as a function of *quantity*, the *demand curve* slopes downward and the *supply curve* slopes upward, as shown in Figure 1(a), since the greater the price of a commodity, the more sellers are inclined to sell and the fewer buyers are willing to buy. Typically, there is some price at which the quantity demanded is equal to the quantity supplied. Graphically, this is the intersection of the demand and supply curves. The price is called the *equilibrium price*, and the corresponding quantity of commodity that is traded is called the *equilibrium quantity*. The equilibrium price and equilibrium quantity are denoted as P_0 and Q_0 respectively in Figure 1(a).

³That is only *intra-marginal* traders are involved in transactions.

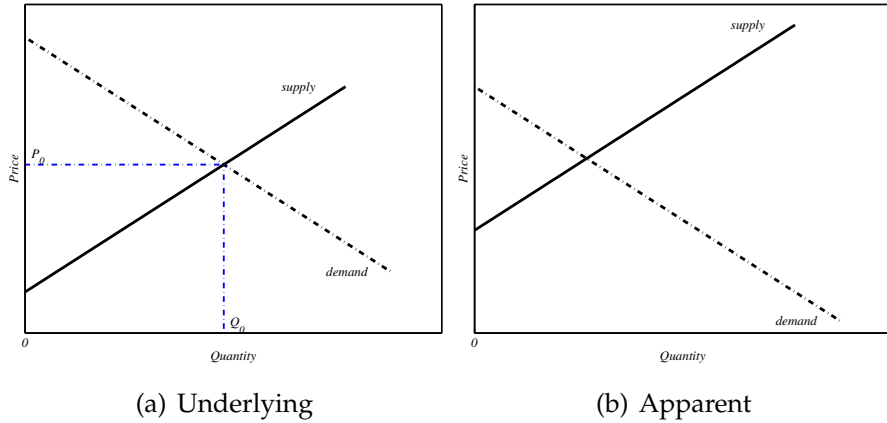


Figure 1: Typical supply and demand curves.

Each trader in an auction presumably has a limit price, called its *private value*, below which sellers will not sell and above which buyers will not buy. The private values of traders are not publicly known in most practical scenarios. What is known instead are the prices that traders offer. Self-interested sellers will presumably offer higher prices than their private values to make a profit and self-interested buyers tend to offer lower prices than their private values to save money. The prices and quantities that are offered also make a set of supply and demand curves, called the *apparent supply and demand curves*, while the curves based on traders' private values are called the *underlying supply and demand*.⁴ Figure 1(b) shows that the apparent supply curve shifts up compared to the underlying supply curve in Figure 1(a), while the apparent demand curve shifts down.

When traders are excessively greedy, the apparent supply and demand curves do not intersect and thus no transactions can be made between sellers and buyers unless they compromise on their profit levels and adjust their offered prices.

1.3 A typical time series of shouts

In a CDA, buyers and sellers not only 'haggle' on prices in a collective manner, but they also face competition from opponents on the same side of the market. Thus buyers, for example, are not only collectively trying to drive prices down, against the wishes of sellers, but they are also individually trying to ensure that they, rather than other buyers, make profitable trades. This leads to shouts becoming more and more competitive over time in a given market. Figure 2 shows a typical time series of shouts

⁴Following the terminology in [7].

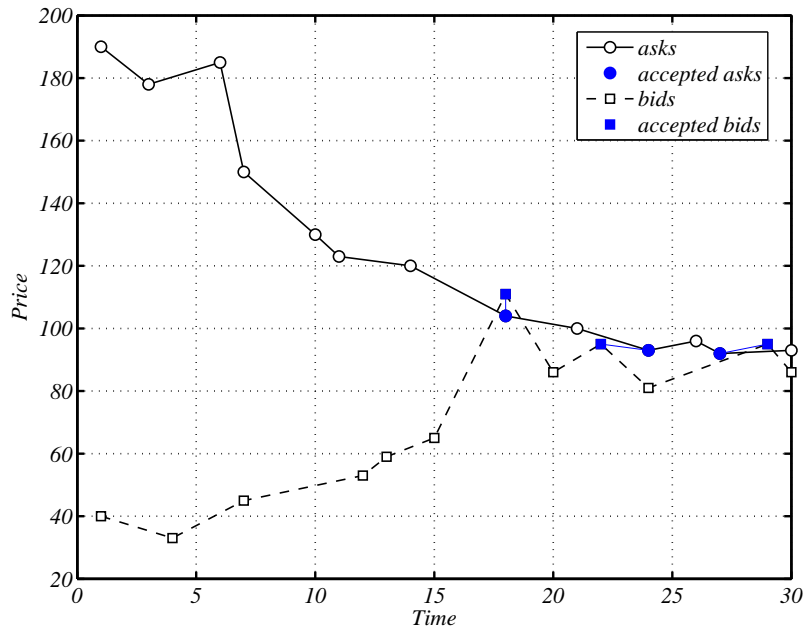


Figure 2: Time series of asks and bids.

in a DA. Ask prices usually start high while bid prices start low. Gradually, traders adjust their offered prices, or make new shouts, closing the gap between standing asks and bids until the price of a bid surpasses that of an ask. Such an overlap results in a transaction, shown as a solid bar between the matched ask and bid in Figure 2.

In the market depicted in Figure 2, newly placed bids (asks) do not have to beat the outstanding bids (asks). However in some variants of the CDA including the market operated by the NYSE, new shouts must improve on existing ones. This requirement is commonly referred to as the NYSE *shout improvement rule* [13].

In some real-world stock markets, including the NYSE and the NASDAQ markets, trades are made through *specialists* or *market makers*, who buy or sell stock from their own inventory to keep the market liquid or to prevent rapid price changes.⁵ Each specialist is required to publish on a regular and continuous basis both a *bid quote*, the highest price it will pay a trader to purchase securities, and an *ask quote*, the lowest price it will accept from a trader to sell securities. The specialist is obligated to stand ready to buy at the bid quote or sell at the ask quote up to a certain number of shares. The range between the lower bid quote and the higher ask quote is called the *bid-ask spread*, which, according to stock exchange regulations, must be suitably small. If buy

⁵In the NYSE, a given stock is traded through a single specialist, and in the NASDAQ, a stock may be dealt with by multiple competing market makers.

orders temporarily outpace sell orders, or conversely if sell orders outpace buy orders, the specialist is required to use its own capital to minimize the imbalance. This is done by buying or selling against the trend of the market until a price is reached at which public supply and demand are once again in balance. Maintaining a bid-ask spread creates risk for a specialist, but when well maintained, also brings huge profits, especially in an active market [1].

Markets involving specialists that present quotes are called *quote-driven markets*. Another class of markets are *order-driven markets*, in which all of the orders of buyers and sellers are displayed. This contrasts with quote-driven markets where only the orders of market makers are shown. An example of an order-driven market is the market formed by *electronic communication networks* or ECNs. These are electronic systems connecting individual traders so that they can trade directly between themselves without having to go through a middleman like a market maker. The biggest advantage of this market type is its transparency. The drawback is that in an order-driven market, there is no guarantee of order execution, meaning that a trader has no guarantee of making a trade at a given price, while it is guaranteed in a quote-driven market. There are markets that combine attributes from quote- and order-driven markets to form hybrid systems.

Our discussion above may give the impression that in real markets trade orders are made directly by the individuals who want to buy or sell stock. In practice, traders commonly place orders through brokerage firms, which then manage the process of executing the orders through a market.⁶

1.4 Performance metrics

Auctions with different rules and populated by different sets of traders may vary greatly in performance. Popular performance measurements include, but are not limited to, *allocative efficiency* and the *coefficient of convergence*.

1.4.1 Allocative efficiency

The allocative efficiency of an auction, denoted as E_a , is used to measure how much social welfare is obtained through the auction. The *actual overall profit*, P_a , of an auc-

⁶<http://www.sec.gov/investor/pubs/tradexec.htm> gives a detailed illustration of how a trade order is executed through a brokerage firm.

tion is:

$$P_a = \sum_i |v_i - p_i| \quad (1)$$

where p_i is the transaction price of a trade completed by agent i and v_i is the private value of agent i , where i ranges over all agents who trade. The *theoretical* or *equilibrium profit*, P_e , of an auction is:

$$P_e = \sum_i |v_i - p_0| \quad (2)$$

for all agents whose private value is no less competitive than the equilibrium price, where p_0 is the equilibrium price. Given these:

$$E_a = \frac{100P_a}{P_e} \quad (3)$$

E_a is thus a measure of the proportion of the theoretical profit that is achieved in practice.

1.4.2 Convergence coefficient

The convergence coefficient, denoted as α , was introduced by Smith [45] to measure how far an active auction is away from the equilibrium point. It actually measures the relative RMS deviation of transaction prices from the equilibrium price:

$$\alpha = \frac{100}{p_0} \sqrt{\frac{1}{n} \sum_{i=1}^n (p_i - p_0)^2} \quad (4)$$

Since markets with human traders often trade close to the equilibrium price, α is used as a way of telling how closely artificial traders approach human trading performance.

2 Game theory

Research on auctions originally interested mathematical economists. They view auctions as games and have successfully applied traditional analytic methods from game theory. This section therefore takes an overlook at basic concepts in game theory.

2.1 Games

The games studied by game theory are well-defined mathematical objects. A game is usually represented in its *normal form*, or *strategic form*, which is a tuple

$$\langle \mathcal{A}_1, \dots, \mathcal{A}_n, R_1, \dots, R_n \rangle.$$

n is the number of players, \mathcal{A}_i is the set of *actions* available to player i , and R_i is the *payoff* or *utility function* $\mathcal{A} \rightarrow \mathfrak{R}$, where \mathcal{A} is the joint action space $\mathcal{A}_1 \times \dots \times \mathcal{A}_n$.

When a player needs to act, it may follow a *pure strategy*, choosing an action, a_i , from its action set, or a *mixed strategy*, π_i , choosing actions according to a probability distribution. The strategy set of player i , denoted as Π_i , is the same thing as a set of probability distributions over \mathcal{A}_i , denoted as $\Delta(\mathcal{A}_i)$. A joint strategy for all players is called a *strategy profile*, denoted as π , and $\pi(a)$ is the probability all players choose the joint action a from \mathcal{A} . Thus player i 's payoff for the strategy profile π is:

$$R_i(\pi) = \sum_{a \in \mathcal{A}} \pi(a) R_i(a).$$

In addition, Π denotes the set of all possible strategy profiles,⁷ π_{-i} is a strategy profile for all players except i , and $\langle \pi_i, \pi_{-i} \rangle$ is the strategy profile where player i uses strategy π_i and the others use π_{-i} .

A normal-form game is typically illustrated as a matrix with each dimension listing the choices of one player and each cell containing the payoffs of players for the corresponding joint action. Figure 3 shows the normal form of the well-known Prisoner's Dilemma game. Alternatively, games may be represented in *extensive form*, which is a tree, as in Figure 4. The tree starts with an initial node and each node represents a state during play. At each non-terminal node, a given player has the choice of action. Different choices lead to different child nodes, until a terminal node is reached where the game is complete and the payoffs to players are given.

⁷It can also be represented as $\Delta(\mathcal{A}_1) \times \dots \times \Delta(\mathcal{A}_n)$.

	cooperate	defect
cooperate	3, 3	0, 4
defect	4, 0	1, 1

Figure 3: The payoff matrix of the Prisoner’s Dilemma game.




Figure 4: A game represented in extensive form.

A game may be *cooperative* or *noncooperative*, as players in these games are respectively *cooperative* or *self-interested*. Cooperative players share a common payoff function, i.e.,

$$\forall i, j \quad R_i = R_j,$$

whereas self-interested players typically have distinct payoff functions. In both cases, players need to coordinate in a certain way to ‘assist’ each other in achieving their goals⁸ [21]. If the payoffs of all players for each strategy profile sum to zero, the noncooperative game is called a *zero-sum* game, i.e.,

$$\forall \pi \in \Pi, \quad \sum_{i=1, \dots, n} R_i(\pi) = 0.$$

Zero-sum games are a special case of a more general class of games called *constant-sum* games, where the sum of all payoffs for each outcome is a constant but may not necessarily be zero. Non-zero-sum games are sometimes referred to as *general-sum* games. In economic situations, the exchange of commodities is considered general-sum, since both parties gain more through the transaction than if they had not transacted (otherwise the exchange would not have happened, assuming both are rational).

⁸The goal of the game designer is also an issue in many situations. Therefore in some parts of the literature, games are considered cooperative as long as they produce a desired systematic outcome even with self-interested players.

In some games, the payoffs for playing a particular strategy remain unchanged as long as the other strategies employed collectively by the players are same, no matter which player takes which action. These games are called *symmetric games* and the rest are *asymmetric games*. For example, the Prisoner's Dilemma game given above is symmetric.

Players may take actions *simultaneously* or *sequentially*. In a sequential game, players have alternating turns to take actions and a player has knowledge about what actions the other players have taken previously. Simultaneous games are usually represented in normal form, and sequential games are usually represented in extensive form. A sequential game is considered a game of *perfect information* if all players know all the actions previously taken by the other players. A similar concept is a game of *complete information*, which means all players in the game know the strategies and payoff functions of the other players. In some sense, complete information may be viewed as capturing static information about a game while perfect information addresses dynamic information that becomes available during runs (or instances) of the game.

2.2 Nash equilibrium

There are various solutions to a normal-form game depending upon the properties of the game and preferences over outcomes.

A strategy, π_i , is said to be *dominant* if it always results in higher payoffs than any other choice no matter what the opponents do, i.e.,

$$\forall \pi'_i \in \Pi, \quad R_i(\langle \pi_i, \pi_{-i} \rangle) \geq R_i(\langle \pi'_i, \pi_{-i} \rangle).$$

In the example of the Prisoner's Dilemma game, the choice **defect** dominates **cooperate** for either player, though ironically both will be better off if they choose to cooperate and know that the other will also.

In many games there are, however, no dominant strategies. To conservatively guarantee the best worst-case outcomes, a player may play the *minimax* strategy, which is

$$\arg \max_{\pi_i \in \Pi_i} \min_{a_{-i} \in \mathcal{A}_{-i}} R_i(\langle \pi_i, \pi_{-i} \rangle),$$

where a_{-i} and \mathcal{A}_{-i} are respectively a joint action for all players except i and the set of joint actions for them. In theory, this can be solved via linear programming, but clearly there are many games that are too large to be solved in practice.

Another approach to solving the problem is to find the *best response* strategies to the strategies of the other players. These can be defined as

$$BR_i(\pi_{-i}) = \{\pi_i | \forall \pi'_i, R_i(\langle \pi_i, \pi_{-i} \rangle) \geq R_i(\langle \pi'_i, \pi_{-i} \rangle)\}.$$

A joint strategy forms a *Nash equilibrium* or NE if each individual strategy is the best response to the others' strategies. When a NE is reached, no player can be better off unilaterally, given that the other players stay with their strategies. In the example of the Prisoner's Dilemma game, $\langle \mathbf{defect}, \mathbf{defect} \rangle$ is a NE.

Although Nash [28] showed that all finite normal-form games have at least one NE, Nash equilibria are generally difficult to achieve. On the one hand, Conitzer and Sandholm [9] proved that computing Nash equilibria is likely NP-hard; on the other hand, some games involve more than one NE, thus without some extra coordination mechanism, no player knows which equilibrium the others would choose.

Many papers have been concerned with "equilibrium refinements" so as to make one equilibrium more plausible than another, however it seems to lead to overly complicated models that are difficult to solve. A more practical approach is to allow players to learn by playing a game repeatedly. A *repeated game* is a game made up from iterations of a single normal-form game, in which a player's strategy depends upon not only the one-time payoffs of different actions but also the history of actions taken by its opponents in preceding rounds. Such a game can be viewed as a system with multiple players and a single state, since the game setting does not change across iterations. If the setting changes over time, the game becomes a *stochastic game*. A stochastic game involves multiple states and the player payoff functions relate to both their actions in each interaction and the current state. The goal of a player in such a game is to maximize its long-term return, which is sometimes defined as the average of all one-time payoffs or the discounted sum of those payoffs.

Brown [2] introduced a learning method, called *fictitious play*, for games in which all the other players use stationary strategies. With this method, the player in question keeps a record of how many times the other players have taken each action and uses the frequencies of actions to estimate the probabilities of actions in an opponent's strategy. Then the player chooses a best-response strategy based on its belief. If the player's belief converges, what it converges to and its own best-response strategy form a NE [16]. The method becomes flawed if players adopt non-stationary strategies and in some games the belief simply does not converge [52].

Another promising approach is to analyze the situation with evolutionary methods. This approach assumes there is a large population of individuals and each strategy

is played by a certain fraction of these individuals. Then, given the distribution of strategies, individuals with better average payoffs will be more successful than others, so that their proportion in the population will increase over time. This, in turn, may affect which strategies are better than others. In many cases, the dynamic process will move to an equilibrium. The final result, which of possibly many equilibria the system achieves, will depend on the initial distribution. The evolutionary, population-dynamic view of games is useful because it does not require the assumption that all players are sophisticated and think the others are also rational, an assumption that is often unrealistic. Instead, the notion of *rationality* is replaced with the much weaker concept of *reproductive success*.

A related concept, considering the overall outcome rather than individual payoffs, is *Pareto optimality*. A strategy profile, π^* , is *Pareto optimal*, or *Pareto efficient*, if there exists no other strategy profile producing higher payoffs for all players, i.e.,

$$\forall \pi \in \Pi, \exists i (R_i(\pi) > R_i(\pi^*)) \Rightarrow \exists j (R_j(\pi^*) > R_j(\pi)).$$

In the Prisoner's Dilemma game, all pure strategy profiles except for **(defect, defect)** are Pareto optimal. A Pareto optimal outcome is highly desirable, but usually difficult to achieve. Self-interested players tend to take locally optimal actions that may not collectively be Pareto optimal. In the Prisoner's Dilemma game, **(cooperate, cooperate)** instead of the NE **(defect, defect)** obviously causes both players to be better off.

In games of incomplete information, to utilize the concept of NE, each player needs to maintain an estimate of the others' strategies so as to come up with a best-response strategy, where Bayes' theorem is used to update or revise beliefs following interactions with opponents. The concept of equilibrium therefore becomes *Bayesian Nash Equilibrium*, or BNE. That is each player's strategy is a function of her own information, and maximizes her expected payoff given other players' strategies and given her beliefs about other players' information [20, 52].

3 Auction theory

Auctions are a way to enable interactions among traders, and traders make profits as a result of transactions. Vickrey [47] pioneered the approach of thinking about a market institution as a game of incomplete information since traders do not know each others' private values.

In research on single-sided auctions, the main goal is to find mechanisms that maximize the profit of sellers,⁹ who are special players in the auctioning games. while in double-sided auctions, research focuses on maximizing social welfare and identifying how price formation develops dynamically.

3.1 Revenue equivalence theorem

By assuming a fixed number of "symmetric",¹⁰ risk-neutral¹¹ bidders, who each want a single unit of goods, have a private value for the object, and bid independently, Vickrey showed that the seller can expect equal profits on average from all the standard types of auctions. This finding is called the *Revenue Equivalence Theorem*. This theorem provides the foundation for the analysis of *optimal auctions*¹² and much subsequent research can be understood in terms of this theorem. Numerous articles have reported how its results are affected by relaxing the assumptions behind it.

The assumption that each trader knows the value of the goods being traded, and that these values are all private and independent of each other is commonly called the *private-value model* [20, 32].

In some cases, by contrast, the actual value of the goods is the same for everyone, but bidders have different private information about what that value actually is. In these

⁹There is no formal distinction between *normal auctions*, in which the auctioneer is the seller and the participants are buyers who have values for the object(s) to be sold, and *procurement auctions*, where the auctioneer is a buyer and the participants are sellers who have costs of supplying the object(s) to be bought.

¹⁰That is bidders' private values are drawn from a common distribution.

¹¹In economics, the term *risk neutral* is used to describe an individual who cares only about the expected return of an action, and not the *risk* (variance of outcomes or the potential gains or losses). A risk-neutral person will neither pay to avoid risk nor actively take risks. Similarly, there are *risk-averse* and *risk-seeking* individuals; they respectively favor the (usually lower) outcome with more certainty and the highest possible outcome (usually with lower probability) [52].

¹²Auctions that maximize the expected profit of sellers.

cases, a bidder will change her estimate of the value if she learns another bidder's estimate, in contrast to the private-value case in which her value would be unaffected by learning any other bidder's preferences or information. This is called the *pure common-value model* [54]. The winner in this scenario is the individual who makes the highest estimate of the value, and this tends to be an overestimate of the value. This overestimation is called the *winner's curse*. If all the bidders have the existence of the winner's curse in mind, the highest bid in first-price auctions tends to be lower than in those second-price auctions, though it still holds that the four standard auctions are revenue-equivalent.

A general model encompassing both the private-value model and the pure common-value model as special cases is the *correlated-value model* [25]. This assumes that each bidder receives a private information signal, but allows each bidder's value to be a general function of *all* the signals.¹³ Milgrom and Weber analyzed auctions in which bidders have *affiliated information*,¹⁴ and showed that the most profitable standard auction is then the ascending auction.

Myerson [26] demonstrated how to derive optimal auctions when the assumption of symmetry fails. Maskin and Riley [23] considered the case of risk-averse bidders, in which case the first-price sealed-bid auction is the most profitable of the standard auctions.

For practical reasons, it is more important to remove the assumptions that the number of bidders is unaffected by the auction design, and that the bidders necessarily bid independently of each other. According to [20], sealed-bid designs frequently (but not always) both attract a large number of serious bidders and are better at discouraging collusion than English auctions.

3.2 On double-sided auctions

In contrast with simple single-sided auctions, where the goals of auction mechanism designers reflect the interests of the single seller, double-sided auctions aim to maximize the collective interests of all traders, or in other words, the social welfare, i.e., the

¹³That is, bidder i receives signal t_i and would have value $v_i(t_1, \dots, t_n)$ if all bidders' signals were available to her. In the private-value model $v_i(t_1, \dots, t_n)$ is a function only of t_i . In the pure common-value model $v_i(t_1, \dots, t_n) = v_j(t_1, \dots, t_n)$ for all i and j .

¹⁴Roughly speaking, bidders' information is affiliated if when one bidder has more optimistic information about the value of the prize, then it is more likely that other bidders' information will also be optimistic.

total surplus all traders earn in an auction. Numerous publications have reported theoretical assertions or empirical observations of high efficiency in a variety of double-sided auctions, and have discussed what leads to the maximization of social welfare.

Chatterjee and Samuelson [4] made the first attempt to analyze double auctions considering a special case of the CH involving a single buyer and a single seller. In this auction, the transaction price is set at the midpoint of the interval of market-clearing prices when the interval is non-empty. They found linear BNE bidding strategies which miss potential transactions with probability of $1/6$. Satterthwaite and Williams [44] analyzed a generalized version of this auction—so-called *k-double auction* or *k-DA*—which involves m sellers and m buyers, and sets the transaction price at other points in the interval of market-clearing prices. They showed that in BNEs the differences between buyers' bids and true values are $O(1/m)$ and foregone gains¹⁵ from trade are $O(1/m^2)$, so *ex post*¹⁶ inefficiency vanishes reasonably fast as the market gets larger.

Wilson [53] first studied the generalization of games of incomplete information to CDAs, in particular, CDAs in which each agent can trade at most one indivisible unit and, given the bids and asks, the maximum number of feasible trades are made at a price a fraction k of the distance between the lowest and highest feasible market clearing prices. He proposed a strategy for buyers and sellers in which a trader waits for a while before making bids or asks. Then the trader conducts a Dutch auction until an offer from the other side is acceptable. This strategy produces a nearly *ex post* efficient final allocation.

Wurman *et al.* [55] carried out an incentive compatibility analysis on a CH which is assumed to have M bids and N asks. They showed that the $(M + 1)$ st-price (or N th-lowest-price) clearing policy is incentive compatible for single-unit buyers under the private-value model, as is the M th-price (or $(N + 1)$ st-lowest-price) auction for sellers. The only way to get incentive compatibility for both buyers and sellers is for some party to subsidize the auction. Myerson and Satterthwaite [27] showed that there does not exist any bargaining mechanism that is individually rational, efficient, and Bayesian incentive compatible for both buyers and sellers without any outside subsidies.

As Friedman [14] pointed out, though theoretically it is natural to model DAS as a

¹⁵Foregone gain means the missed profit compared with the profit that would have been made if the market cleared at the equilibrium price.

¹⁶This refers to the value that is actually observed or the value calculated after an event occurs, in contrast to *ex ante*, which means the expected value calculated before the resolution of uncertainty.

game of incomplete information, the assumption of prior common knowledge in the incomplete information approach may not hold in continuous auctions or may involve incredible computational complexity. This is because at every moment, a trader needs to compute expected utility-maximizing shouts based on the shout and transaction history of the auction and the length of time the auction has to go. On the other hand, laboratory results have shown that DA outcomes are quite insensitive to the number of traders beyond a minimal two or three active buyers and two or three active sellers.¹⁷ Moreover, parameter choices, which according to an incomplete information analysis, should greatly reduce efficiency in DAs had no such effect in recent laboratory tests [14].

¹⁷This observation and the above result of $O(1/m^2)$ foregone gains by Satterthwaite and Williams [44] may suggest that the coefficient of $1/m^2$ in the actual foregone gain function is small.

4 Experimental approaches

Due to the difficulty of applying game-theoretic methods to complex auction mechanisms, researchers from economics and computer science have turned to running laboratory experiments, studying the dynamics of price formation and how the surprisingly high efficiency is obtained in a DA where information is scattered between the traders.

4.1 Different data sources

The data on which researchers base their studies may come from three sources: (1) field data from large-scale on-going markets, (2) laboratory data from small-scale auctions with human subjects, and (3) computer simulation experiments.

Field data has the most relevance to the real-world economy, but does not reveal many important values, e.g., the private values of traders, and hence puts limits on what can be done. The human subjects in laboratory experiments presumably inherit the same level of intelligence and incentive to make profit as in real markets, and the experiments are run under the rules that researchers aim to study. Such experiments, however, are expensive in terms of time¹⁸ and money¹⁹ needed.

Computer-aided simulation is a less expensive alternative and can be repeated as many times as needed. However traders' strategies are not endogenously chosen as in auctions with human traders, but are specified exogenously by the experiment designers, which raises the question of whether the conclusions of this approach are trustworthy and applicable to practical situations. Gode and Sunder [18] invented a zero-intelligence strategy (or ZI) that always randomly picks a profitable price to bid or ask. Surprisingly, their experiments with CDAs exhibit high efficiency despite the lack of intelligence of the traders. Thereafter, much more work followed this path, and has gained tremendous momentum, especially given that real-world stock exchanges are becoming automated, e-business becomes an everyday activity, and the Internet reaches every corner of the globe.

¹⁸The experiments are run using a physical clock and need take into consideration the response time of human traders.

¹⁹Usually human subjects are monetarily rewarded according to their performance.

4.2 Smith's experiments

Smith pioneered the research falling into the so-called *experimental economics* field by running a series of experiments with human subjects [45]. The experimental results revealed many of the properties of CDAs, which have been the basis and benchmark for much subsequent work. Smith showed that in many different cases even a handful of traders can lead to high allocative efficiency, and transaction prices can quickly converge to the theoretical equilibrium.

Smith's experiments are set up as follows:

- Every trader, either a buyer or a seller, is given a private value. The set of private values form the supply and demand curves.
- Each experiment was run over a sequence of trading days, or periods,²⁰ the length of which depend on how many traders are involved but are typically several minutes in duration. Different experiments may have different numbers of periods.
- For simplicity, in most experiments, a trader is allowed to make a transaction for the exchange of only a single commodity in each day.
- Traders are free at any time to make a bid/ask or to accept a bid/ask.
- Once a transaction occurs, the transaction price, as well as the two traders' private values, are recorded.
- For each new day, a trader may make up to one transaction with the same private value as before no matter whether she has made one in the previous day. Thus the supply and demand curves each correspond to a trading day. The experimental conditions of supply and demand are held constant over several successive trading days in order to give any equilibrating mechanisms an opportunity to establish an equilibrium over time, unless it is the aim to study the effect of changing conditions on market behavior.

[45] reports 10 experiments that we discuss below. Each experiment was summarized by a diagram showing the series of transactions in the order in which they occurred.

²⁰Smith used the term *periods* to refer to what is called *days* in this report.




Figure 5: Supply and demand curves (left) and transaction price trajectory (right) in Smith's test 1. Originally as Chart 1 in [45].

Figure 5 gives one of Smith's diagrams.²¹ In the right-hand part of the diagram, each tick represents a transaction, rather than a unit of physical time.

Trading prices in most experiments have a striking tendency to converge on the theoretical prices, marked with a dashed line in Figure 5. To measure the tendency to converge, Smith introduced the coefficient of convergence, α , from (4). Figure 5 shows α tends to decline from one trading day to the next.

The equilibrium price and quantity of experiments 2 and 3 are approximately the same, but the latter, with the steeper inclination of supply and demand curves, converges more slowly. This complies with the Walrasian hypothesis that the rate of increase in exchange price is an increasing function of the excess demand at that price.

Experiment 4 presents an extreme case with a flat supply curve, whose result also confirms the Walrasian hypothesis, but it converges to a fairly stable price above the predicted equilibrium. In this experiment, a decrease in demand is ineffective in shocking the market down to the equilibrium. The result shows that the equilibrium

²¹Supply and demand curves of a market are typically stepped due to the discrete numbers of commodities, but the ones in Figure 1 are straight line segments because it is assumed there that a large number of traders participate in the auction and thus the step-changes can be treated as infinitesimal.

may depend not only on the intersection of the supply and demand schedules, but also upon the shapes of the schedules.

A hypothesis aiming to explain the phenomenon is that the actual market equilibrium will be above the equilibrium by an amount which depends upon how large the *buyers' rent*²² is relative to the *sellers' rent*.²³ Experiment 7, which was designed with the purpose of supporting or contradicting this hypothesis, shows slow convergence, complying with the Walrasian hypothesis, but still exhibits a gradual approach to equilibrium. It is concluded that a still smaller buyers' rent may be required to provide any clear downward bias in the static equilibrium. What's more, it seems "quite unmistakable" that the bigger the difference between the buyers' rent and sellers' rent, the slower the convergence. Smith speculated that the lack of monetary payoffs to the experimental traders may have an effect on the markets. A strong measure to further test the hypothesis is to mimic real markets as exactly as possible by paying each trader a small return just for making a contract in any period, which according to some experiments induces faster convergence.

Experiment 5 was designed to study the effect on market behavior of changes in the conditions of demand and supply. At some point in the experiment, new buyers were introduced resulting in an increase in demand. The eagerness to buy causes the trading price to increase substantially once the market resumes and the price surpasses the previous equilibrium.

Experiment 6 was designed to determine whether market equilibrium was affected by a marked imbalance between the number of intra-marginal sellers and the number of intra-marginal buyers near the predicted equilibrium price. The result confirmed the effect of a divergence between buyer and seller rent on the approach to equilibrium, but the lack of marginal sellers near the theoretical equilibrium did not prevent the equilibrium from being attained. The change of decrease in demand at the end of the fourth trading day showed that the market responded promptly by showing apparent convergence to the new, lower, equilibrium.

In contrast to the previous experiments, the market in experiment 8 was designed to simulate an ordinary retail market, in which only sellers are allowed to enunciate offers, and buyers could only either accept or reject the offers of sellers. Due to the desire of sellers to sell at higher prices, the trading prices in the first period remained above the predicted equilibrium. But starting at the second period, the trading price

²²The area enclosed by the horizontal line at P_0 , price axis, and the demand curve.

²³The area enclosed by the horizontal line at P_0 , price axis, and the supply curve.

decreased significantly and remained below the equilibrium, not only because the early buyers again refrained from accepting any high price offers, but also because the competition among sellers became more intense. Later in the experiment, when the previous market pricing organization was resumed, exchange prices immediately moved toward equilibrium.

Experiments 9 and 10 are similar to experiment 7 except that each trader is allowed to make up to 2 transactions with the assigned private value within each day. The results showed that the increase in volume helps to speed up the convergence to equilibrium. The same results were obtained even when demand was increased during experiment 9.

5 Trading agents

5.1 Zero intelligence traders

Smith’s focus in [45] was mainly on the convergence of transaction prices in different scenarios rather than directly examining why high efficiency is obtained. However, high efficiency is usually the goal of a DA market designer. In a computerized world, a question that arises naturally is whether Smith’s results can be replicated in electronic auctions. In Smith’s experiments, as is traditional in real markets, the traders are human beings, but computer programs are supposed to be automatic and work without human involvement. Obviously humans are intelligent creatures, but programs are not, at least for the foreseeable future. Is it intelligence that contributes to the high efficiency of double auction markets, or is it something else?

Gode and Sunder [18, 19] were among the first to address this question, claiming that no intelligence is necessary for the goal of achieving high efficiency; so the outcome is due to the auction mechanism itself.

They reached this position having introduced two trading strategies: *zero intelligence without constraint* or ZI-U and *zero intelligence with constraint* or ZI-C. ZI-U, the more naïve version, shouts an offer at a random price without considering whether it is losing money or not, while ZI-C, which lacks the motivation of maximizing profit and picks a price in a similar way to ZI-U, simply makes shouts that guarantee no loss.

It was shown that ZI-U performs poorly in terms of making a profit, but ZI-C generates high efficiency solutions, comparable to the human markets (see Table 1) and can be considered to place a lower bound on the efficiency of markets [19].

Gode and Sunder’s experiments were setup with similar rules as in Smith’s. They designed five different supply and demand schedules and tested each of them respec-

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: Mean efficiency of markets in Gode and Sunder’s experiments. Originally as Table 2 in [18].

tively with the three kinds of homogeneous traders, ZI-U, ZI-C, and human traders. Figure 6 presents what happened in one of their experiments.

Prices in the ZI-U market exhibit little systematic pattern and no tendency to converge toward any specific level, but on the contrary, prices in the human market, after some initial adjustments, settle in the proximity of the equilibrium price (indicated by a solid horizontal line in all panels in Figure 6). Gode and Sunder then raised the question: how much of the difference between the market outcomes with ZI-U traders and those with human traders is attributable to intelligence and profit motivation, and how much is attributable to market discipline?

They argue that, after examining the performance of the ZI-C markets, it is market discipline that plays a major role in achieving high efficiency. Though in the ZI-C market, the price series shows no signs of improving from day to day, and the volatility of the price series is greater than the volatility of the price series from the human market, the series converges slowly toward equilibrium within each day. Gode and Sunder's explanation is that it is due to the progressive narrowing of the opportunity sets of ZI-C traders, e.g., the set of intra-marginal traders. Despite the randomness of ZI-C, buyers with higher private values tend to generate higher offered prices and they are likely to trade with sellers earlier than those buyers further down the demand curve. A similar statement also holds for sellers. Thus as the auction goes on, the upper end of the demand curve shifts down and the lower end of the supply curve moves up, which means the feasible range of transaction prices narrows as more commodities are traded, and transaction prices will converge to the equilibrium price. The fact that ZI-C traders lack profit motivation and have only the minimal intelligence (just enough to avoid losing money) suggests that the market mechanism is the key to obtaining high efficiency.

5.2 Zero intelligence plus and beyond

Gode and Sunder's results were, however, questioned by Cliff and Bruten [7]. The latter agreed on the point that the market mechanism plays a major role in achieving high efficiency, but disputed whether in ZI-C markets transaction prices will always converge on equilibrium price. They argued that the mean or expected value of the transaction price distribution was shown quantitatively to get close to the equilibrium price only in situations where the magnitude of the gradient of linear supply and demand curves is roughly equal, and used this to infer that zero-intelligence traders are not sufficient to account for convergence to equilibrium.

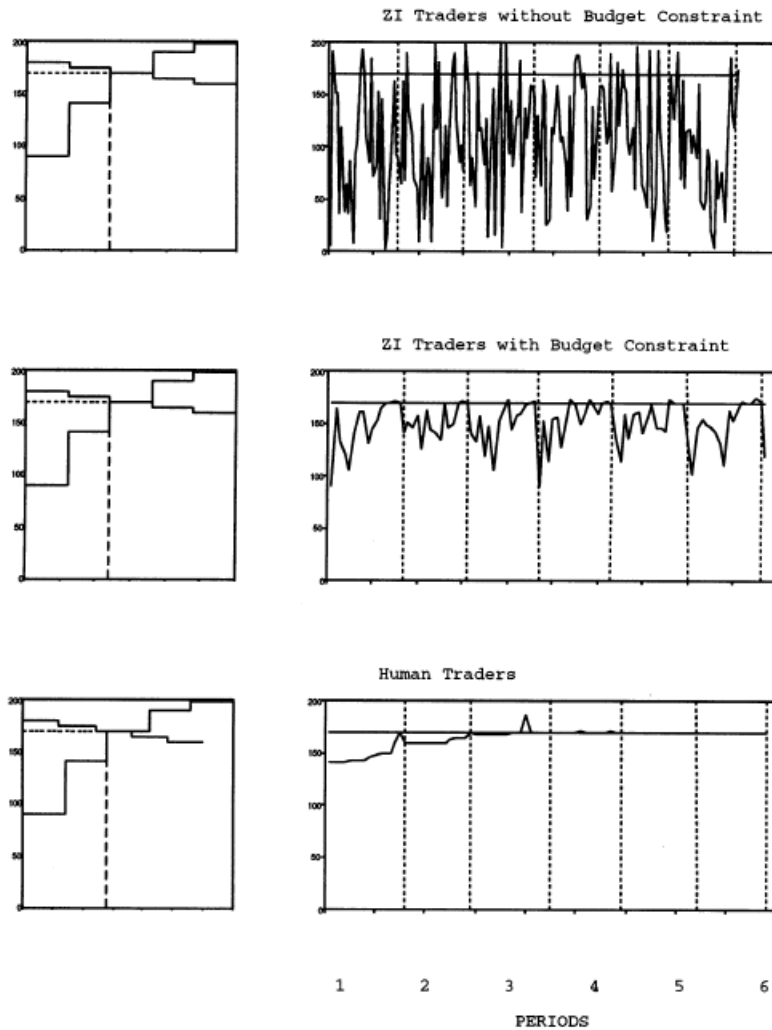


Figure 6: Gode and Sunder's experiments comparing ZI-U traders (top), ZI-C traders (middle), and human traders (bottom). Originally as Fig. 4 in [18].

Cliff and Bruten further designed an adaptive trading strategy called *zero intelligence plus* or ZIP. Like ZI-C, ZIP traders make stochastic bids, but can adjust their prices based on the auction history, i.e., raising or lowering their profit margins dynamically according to the actions of other traders in the market. More specifically, ZIP traders raise the profit margin when a less competitive offer from the competition²⁴ is accepted, and lower the profit margin when a more competitive offer from the competition is rejected, or an accepted offer from the other side of the market would have been rejected by the subject. At every step, the profit margin is updated according to

²⁴That is sellers compete against sellers to get asks accepted and buyers compete against buyers to get bids accepted.

a learning algorithm called the *Widrow-Hoff delta rule* in which a value being learned is adapted gradually towards a moving target, and the past targets leave discounting momentum to some extent.

Cliff and Bruten concluded that the performance of ZIP traders in the experimental markets is significantly closer to that of human traders than is the performance of ZI-C traders, based on the observation that ZIP traders rapidly adapt to give profit dispersion²⁵ levels that are in some cases approximately a factor of ten less than those of ZI-C traders.

Preist and van Tol introduced a revised version of ZIP, which we call PVT, and reported faster convergence to equilibrium and robustness to changes in parameter configuration [39].

Other learning methods have been adopted to design even more complex trading strategies than ZIP and its variants. Roth and Erev [40] proposed a reinforcement-based stimuli-response strategy, which we call RE. RE traders adapt their trading behavior in successive auction rounds by using their profits in the last round as a reward signal. Gjerstad and Dickhaut [17] suggested a best-response-based strategy, which is commonly referred to as GD. GD traders keep a sliding window of the history of the shouts and transactions and calculate the probabilities of their offers being accepted at different prices. The traders use a cubic interpolation on the shouts and transaction prices in the sliding window in order to compute the probability of future shouts being accepted. They then use this to calculate the expected profit of those shouts. The expected profit at a price is the product of the probability of the price being accepted and the difference between the price and the private value. GD traders then always choose to bid or ask at a price that maximizes their expected profit. GD is the most computation-intensive trading strategy considered so far, and indeed generates the best record both for allocative efficiency and the speed of convergence to equilibrium compared to the other trading strategies in literature.

By way of indicating typical efficiencies achieved in a CDA, Figure 7 shows the trend of the overall efficiencies of homogeneous CDAs lasting 10 days with 50 rounds per day in which 10 buyers and 10 sellers all use the same strategy, one of: TT,²⁶ KA-

²⁵Profit dispersion is the root mean squared difference between actual and equilibrium profits, and can be expressed as

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

where a_i and π_i are the actual and theoretical equilibrium profits of trader i , $i = 1, \dots, n$.

²⁶TT denotes the Truth-telling strategy, in which agents truthfully report their private values.




Figure 7: E_a s for CDA markets populated by common trading strategies




Figure 8: Supply and demand schedules in the CDA markets presented in Figure 7.

PLAN,²⁷ ZIP, RE, and GD. The results are averaged over 400 iterations and obtained in JASA—the extensible Java-based auction simulation environment [33]. Figure 8 gives the supply and demand schedules in the markets.

²⁷‘Kaplan’ refers to Todd Kaplan’s sniping strategy, in which agents wait until the last minute before attempting to steal the deal [42]. See Section 5.3 for more information.

5.3 Interaction of heterogenous trading strategies

All the above empirical works have employed either human traders or homogeneous trading agents, demonstrating high efficiency and fast convergence to equilibrium, and some of this work has also produced theoretical results. It is however necessary to see how an auction works populated by heterogeneous trading subjects.

There are both theoretical and practical reasons for considering heterogeneous traders. As Rust *et al.* argued in [42]:

Although current theories of DA markets have provided important insight into the nature of trading strategies and price formation, it is fair to say that none of them has provided a satisfactory resolution of “Hayek’s problem”.²⁸ In particular, current theories assume a substantial degree of implicit coordination by requiring that traders have common knowledge of each other’s strategies (in game-theoretic models), or by assuming that all traders use the same strategy (in learning models). Little is known theoretically about price formation in DA markets populated by heterogeneous traders with limited knowledge of their opponents.

... the assumption that players have common knowledge of each other’s beliefs and strategies ... presumes an unreasonably high degree of implicit coordination amongst the traders ... Game theory also assumes that there is no *a priori* bound on traders’ ability to compute their BNE strategies. However, even traders with infinite, costless computing capabilities may still decide to deviate from their BNE strategies if they believe that limitations of other traders force them to use a sub-optimal strategy.

They went on to argue that ZI-C and other strategies’ striking performance strongly suggests that the nice properties have more to do with the market mechanism itself than the rationality of traders. In addition, strategies that are more individually rational than ZI-C may display less collective rationality since clever strategies can exploit unsophisticated ones such as TT and ZI-C so that a more-intelligent extra-marginal trader has more chances to finagle a transaction with an intra-marginal traders, causing market efficiency to fall.

To observe heterogeneous auctions, the Santa Fe Double Auction Tournament (SFDAT) was held in 1990 and prizes were offered to entrants in proportion to the trading

²⁸That is how the trading process aggregates traders’ dispersed information, driving the market towards competitive equilibrium.

profits earned by their programs over the course of the tournament. 30 programs from researchers in various fields and industry participated. The majority of the programs encoded the entrant's "market intuition" using simple rules of thumb. The top-ranked program was KAPLAN, named after the entrant. KAPLAN and the runner-up strategy are remarkably similar. Both "wait in the background and let the others do the negotiating, but when bid and ask get sufficiently close, jump in and steal the deal" [42].

The overall efficiency levels in the markets used in the tournaments originally appear to be somewhat lower than that observed in experimental markets with human traders, but experiments without the last-placed players produced an efficiency of around 97%. This is further evidence that the properties of traders also affect the outcome of DA markets to some extent.

Besides high efficiency levels and convergence to competitive equilibrium, other "stylized facts" of human DA markets observed in the SFDAT include: reductions in transaction-price volatility and efficiency losses in successive trading days that seem to reflect apparent learning effects, coexistence of extra-marginal and intra-marginal efficiency losses, and low-rank correlations between the *realized order of transactions* and the *efficient order*.²⁹

Thorough examination of efficiency losses in the tournaments and later experiments indicates that the success of KAPLAN is due to its patience in waiting to exploit the intelligence or stupidity of other trading strategies.³⁰

The volume of e-commerce nowadays creates another motivation for evaluating trading strategies in a heterogeneous environment. Electronic agents, on behalf of their human owners, can automatically make strategic decisions and respond quickly to the changes in various kinds of markets. In the foreseeable future, these agents will have to compete with a variety of agents using a range of trading strategies and human traders. As more complex trading strategies appear, it is natural to speculate on how these electronic minds will compete against their human counterparts.

Das *et al.* ran a series of CDAs allowing persistent orders³¹ populated by a mixed

²⁹The *efficient order* is the transaction sequence that maximizes surplus, meaning that the first transaction occurs between the buyer with the highest private value and the seller with the lowest private value, the second transaction occurs between the buyer and seller next to them, and so on. The *realized order of transactions* is the actual order in which transactions are made.

³⁰The usual higher efficiency of CHs than CDAs can also be viewed as the proactive elimination of the effect of traders' impatience.

³¹In the SFDAT and the CDA testing ZIP in [7], shouts that are outbid are removed from the market, which is however not typical of real marketplaces.

population of automated agents (using modified GD and ZIP strategies) and human traders [12]. They found that though the efficiency of the CDAs was comparable with prior research, the agents outperformed the humans in all the experiments, obtaining about 20% more profit. Das *et al.* speculated that this was due to human errors or weakness, and human traders were observed to improve their performance as they got familiar with using the trading software. Das *et al.* also suggested that the weaknesses of trading agents may be found when human experts take them on and thus improvement can be made to the algorithms of the trading agents.³²

Tesauro and Das [46] executed experiments with both homogeneous and heterogeneous trading agents with varying trader population composition, making it possible to gain more insights into the relative competitiveness of trading strategies. In either the so-called “one-in-many”³³ tests or “balanced-group”³⁴ tests, GD and ZIP (and their variants) exhibited superior performance over ZI-C and KAPLAN even when the market mechanisms vary to some extent.³⁵ Furthermore, MGD, a variant of GD due to Das *et al.* [12], outperformed all the other strategies.

The above approaches nevertheless all employ a fixed competition environment. In practice, when a strategy dominates others, it tends to flourish and be adopted by more people. Rust *et al.* are the first that we are aware of to conduct evolutionary experiments, where the relative numbers of the different trading strategies changed over time, so that more profitable strategies became more numerous than less profitable ones. Such an analysis revealed that although KAPLAN agents outperformed others when traders of different types are approximately evenly distributed, they later exhibited low overall efficiency as they became the majority, making the evolution process a cycle of ups and downs.

Walsh *et al.* [49] gave a more formal analysis combining the game-theoretic solution concept of NE and *replicator dynamics*. They treated heuristic strategies, rather than the atomic actions like a bid or ask, as primitive, and computed expected payoffs of each individual strategy at certain points of the joint heuristic strategy space.³⁶

³²[12] also reported that either buyers consistently exploited sellers, or vice versa. However no convincing analysis was given. Similar phenomenon also occurred in experiments described in [38]. It is not clear whether this is caused by the inherent randomness in the trading agents.

³³A single agent of one type competes against an otherwise homogeneous population of a different type.

³⁴Buyers and sellers are evenly split between two types, and every agent of one type has a counterpart of the other type with identical limit prices.

³⁵[46] tested both with and without the NYSE shout improvement rule, a standing shout queue, and allowance of shout modification.

³⁶That is a space of a mixture of strategies when their relative proportions vary.

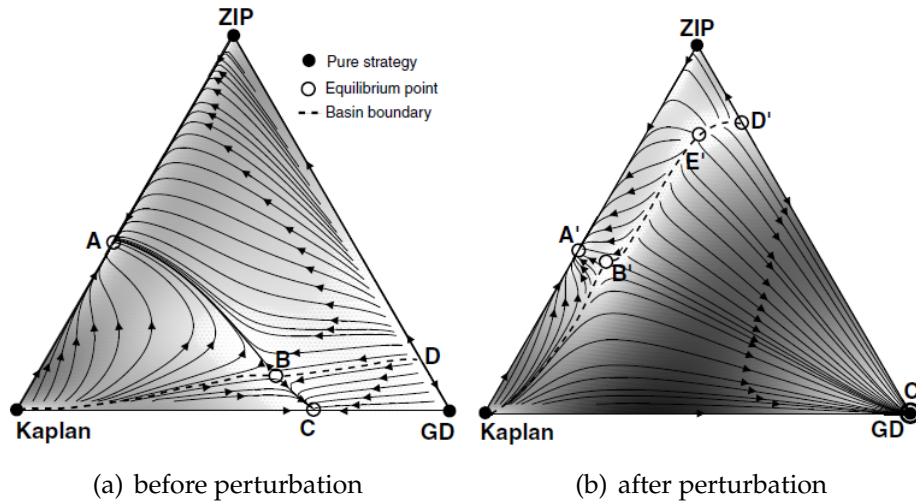


Figure 9: The replicator dynamics of CDA with ZIP, KAPLAN, and GD. Originally as Figure 2 in [49].

This method reduced the model of the game from a potentially very complex, multi-stage game to a one-shot game in normal form. At points where one strategy gains more than others, replicator dynamics dictates that the whole population moves to a nearby point where the winning strategy takes a larger fraction of the population. This process continues until an equilibrium point is reached where either the population becomes homogeneous or all strategies are equally competitive in terms of their expected payoffs. There may be multiple equilibrium points³⁷ ‘absorbing’ areas of different sizes, *basins* of the equilibria, which together compose the whole strategy space. In particular, Figure 9(a) shows the replicator dynamics of a CDA market with three strategies. A , B , C , and D are all equilibrium points, but B and D are not stable since a small deviation from them will lead to one of the other equilibria. The triangle field gives an overview of the interaction of the three strategies and their relative competitiveness. What’s more, a technique called *perturbation analysis* is used to evaluate the potential to improve on a strategy. Figure 9(b) shows the replicator dynamics of the same strategies after small portions of both ZIP and KAPLAN’s payoffs were shifted to GD. Such a shift significantly changed the landscape of the space, and GD dominated in most of possible combinations. This showed that a ‘tiny’ improvement on the GD strategy may greatly affect its competition against the other strategies.

Phelps *et al.* [35, 34] took a similar approach in comparing the RE, TT, and GD strategies, showed the potential of RE, and demonstrated that a modified RE strategy could

³⁷Each equilibrium point also represents a mixed strategy, a homogeneous population of which makes a NE.

be evolved by optimizing its learning component.

The main drawback of this approach is an exponential dependence on the number of strategies, which limits its applicability to real-world domains where there are potentially many heuristic strategies. Walsh *et al.* [50] proposed information theoretic approaches to deliberately choose the sample points in the strategy space through an interleaving of equilibrium calculations and payoff refinement, thus reducing the number of samples required.

5.4 Automating strategy acquisition

Designing heuristic strategies to a great extent depends on the intelligence and experience of the strategy designer. Prior studies have also demonstrated that heuristic strategies' performance hinges on the selection of parameter values. Automatic optimization is preferable in this sense to find best parameter combinations and further identify better strategies. Cliff and Phelps *et al.* are the pioneers in this work, adopting evolutionary computation to address the challenge.

5.4.1 Evolutionary computation

A *genetic algorithm* (or GA) is a search technique used in computing to find true or approximate solutions to optimization and search problems. Genetic algorithms are a particular class of evolutionary algorithms that use techniques inspired by evolutionary biology such as *inheritance*, *mutation*, *selection*, and *crossover* [52].

A typical genetic algorithm requires two things to be defined:

1. a *genetic representation* of the solution domain, also called the *genotype* or *chromosome* of the solution species,
2. a *fitness function* to evaluate the solution domain.

A standard representation of the solution is as an array of bits. Arrays of other types and structures can be used in essentially the same way. The main property that makes these genetic representations convenient is that their parts are easily aligned due to their fixed size, which facilitates simple crossover operation. Variable length representations have also been used, but crossover implementation is more complex in this case.

The fitness function is defined over the genetic representation of a solution and measures the quality of the solution. The fitness function is always problem dependent. For instance, in the knapsack problem we want to maximize the total value of objects that we can put in a knapsack of some fixed capacity. A representation of a solution might be an array of bits, where each bit represents a different object, and the value of the bit (0 or 1) represents whether or not the object is in the knapsack. Not every such representation is valid, as the size of objects may exceed the capacity of the knapsack. The fitness of the solution is the sum of values of all objects in the knapsack if the representation is valid, or 0 otherwise. In some problems, it is hard or even impossible to define the fitness expression; in these cases, interactive genetic algorithms are used.

Once we have the genetic representation and the fitness function defined, the GA proceeds to initialize a population of solutions randomly, then improve it through repetitive application of mutation, crossover, and selection operators.

5.4.2 Optimizing parameter combination in ZIP

Cliff addressed the labor-intensive manual parameter optimization for the ZIP strategy, automatically optimizing parameter selection using a GA [6]. He identified eight parameters in ZIP: lower and upper bounds of the learning rate β (how fast to move towards the target), momentum γ (how much past momentum to carry over), and initial profit margin μ , and the upper bounds of the ranges defining the distributions of absolute and relative perturbations on learned prices, respectively denoted as c_a and c_r . These real parameters make an eight-dimensional space and any parameter value combination corresponds to a point in that space. The vector of the eight parameters defines an ideal genotype.

5.4.3 Combining GA and heuristic strategy analysis

Phelps *et al.* took a step further along this track. They combined the replicator-dynamics-based heuristic strategy analysis method in [49] and a GA, identified a strategy as the basis for optimization, and successfully evolved the strategy and acquired an optimized strategy that can beat GD, commonly considered the most competitive strategy [35, 34].

Since it is not realistic to seek “best”, or even “good”, strategies that can beat all potential opponents because an absolutely dominating strategy does not appear to exist

in the CDA trading scenario—since the performance of a strategy depends greatly on the types of the opponents—Phelps *et al.* proposed using a small finite population of randomly sampled strategies to approximate the game with an infinite strategy population consisting of a mixture of all possible strategies. In particular, RE, TT, and GD were chosen as sample strategies. Following the heuristic strategy analysis and perturbation method in [49], RE was found to have the potential to dominate TT and GD.

The RE strategy uses reinforcement learning to choose from n possible profit margins over the agent’s private value based on a reward signal computed as a function of profits earned in the previous round of bidding. Potentially, the RE learning algorithm may be replaced by a number of learning algorithms, including SQ (stateless Q-learning), NPT (a modified version of RE used in [29]), and DR (a control algorithm which selects a uniformly random action regardless of reward signal). Phelps *et al.* then encoded the genotype to select any of these algorithms together with their parameters. The evolutionary search procedure they used is similar to Cliff’s except that the individuals in a generation are evaluated again with the heuristic strategy analysis approach and the basin size is used as a measure of fitness. The experiment finally found a SQ algorithm with a particular parameter combination, which together with TT composes the Nash equilibrium that captures 97% of the strategy space populated by the learned strategy, TT, RE, and GD.

5.5 Trading Agent Competition

The Trading Agent Competition (TAC) was organized to promote and encourage high quality research into trading agents. Under the TAC umbrella, a series of competitions have been held, including two types of game, TAC Classic and TAC SCM [51].

TAC Classic sets up a “travel agent” scenario based on complex procurement in multiple simultaneous auctions. Each travel agent (an entrant to the competition) has the goal of assembling travel packages (from TACTown to Tampa, during a notional multi-day period). Each agent is acting on behalf of a certain number of clients, who express their preferences for various aspects of the trip. The objective of the travel agent is to maximize the total satisfaction of its clients (the sum of the client utilities).

TAC SCM was designed to capture many of the challenges involved in supporting dynamic supply chain practices in the industry of PC manufacturing. Supply chain management is concerned with planning and coordinating the activities of organizations across the supply chain, from raw material procurement to the delivery of

finished goods. In today's global economy, effective supply chain management is vital to the competitiveness of manufacturing enterprises as it directly impacts their ability to meet changing market demands in a timely and cost effective manner. In TAC SCM, agents are simulations of small manufacturers, who must compete with each other for both supplies and customers, and manage inventories and production facilities.

6 Experimental auction mechanism design

Mechanism design applied to auctions explores how to design the rules that govern auctions to obtain specific goals.

The story of trading strategies in the preceding section is only one facet of the research on auctions. Gode and Sunder's results suggest that auction mechanisms play an important role in determining the outcome of an auction, and this is further borne out by the work of Walsh *et al.* [49], which also points out that results hinge on both auction design and the mix of trading strategies used.

According to classical auction theory, if an auction is *strategy-proof* or *incentive compatible*, traders need not bother to conceal their private values and in such auctions complex trading agents are not required. However, typical DAS are not strategy-proof. McAfee [24] has derived a form of double auction that is strategy-proof, though this strategy-proofness comes at the cost of lower efficiency.

Despite the success of analytic approaches to the relatively simple auctions presented in Section 3, the high complexity of the dynamics of some other auction types, especially DAS, makes it difficult to go further in using analytical methods [22, 43, 49].

As a result, researchers turned to empirical approaches using machine learning techniques, sometimes combined with methods from traditional game theory. Instead of trying to design optimal auction mechanisms, the computational approach looks for relatively good auctions and aims to make them better, in a noisy economic environment with traders that are not perfectly rational.

6.1 A parameterized space of auctions

One can think of different forms of auctions as employing variations of a common set of the auction rules, forming a parameterized auction space. Wurman *et al.* and others parameterized auction rules using the following classification [41, 56, 57]:

- Bidding rules: determine the semantic content of messages, the authority to place certain types of bids, and admissibility criteria for submission and withdrawal of bids.
 - How many sellers and buyers are there?

- Are both groups allowed to make shouts?
- How is a shout expressed?
- Does a shout have to beat the corresponding market quote if one exists?
- Information revelation:
 - When and what market quotes are generated and announced?
 - Are shouts visible to all traders?
- Clearing policy:
 - When does clearing a market take place?
 - When does a market close?
 - How are shouts matched?
 - How is a transaction price determined?

The idea of parameterizing auction space not only eases the heuristic auction mechanism design, but also makes it possible to ‘search’ for better mechanisms in an automated manner [8, 36, ?].

It is not yet clear how auction design, and thus the choice of parameter values, contributes to the observed performance of auctions. Thus it is not clear how to create an auction with a particular specification. It *is* possible to design simple mechanisms in a provably correct manner from a specification, as shown by Conitzer and Sandholm [10, 11]. However it is not clear that this kind of approach can be extended to mechanisms as complex as DAS. As a result, it seems that we will have to design double auction mechanisms experimentally, at least for the foreseeable future.

Of course, doing things experimentally does not solve the general problem. A typical experimental approach is to fix all but one parameter, creating a one-dimensional space, and then measure performance across a number of discrete sample points in the space, obtaining a fitness landscape that is expected to show how the factor in question correlates to a certain type of performance and how the auction can be optimized by tweaking the value of that factor [38]. In other words, the experimental approach examines one small part of a mechanism and tries to optimize that part.³⁸ The situation is complicated when more than one factor needs to be taken into consideration — the search space then becomes complex and multiple dimensional, and the computation required to map and search it quickly becomes prohibitive.

³⁸And of course there are rarely any guarantees as to the optimality of the results.

6.2 Evolving market mechanisms

Instead of manual search, some researchers have used evolutionary computation to automate mechanism design in a way that is similar to the evolutionary approach to optimizing trading strategies.

Cliff [5] explored a continuous space of auction mechanisms by varying the probability of the next shout (at any point in time) being made by a seller, denoted by Q_s . The continuum includes the CDA ($Q_s = 0.5$) and also two purely single-sided mechanisms that are similar to the English auction ($Q_s = 0.0$) and the Dutch auction ($Q_s = 1.0$). Cliff's experiments used genetic algorithms and found that a Q_s that corresponds to a completely new kind of auction led to a better α value than that obtained for other markets using ZIP traders. Walia *et al.* [48] and the same authors but in a different order [8] continued with this work, showing that the approach is also effective in markets using ZI-C traders, and the new "irregular" mechanisms can lead to high efficiency with a range of different supply and demand schedules as well. The visualization of fitness landscapes, using plots including 3D histograms and contours, is also noteworthy.

Byde [3] took a similar approach in studying the space of auction mechanisms between the first and second-price sealed-bid auctions. The winner's payment is determined as a weighted average of the two highest bids, with the weighting determined by the auction parameter. For a given population of bidders, the revenue-maximizing parameter is approximated by considering a number of parameter choices over the allowed range, using a GA to learn the parameters of the bidders' strategies for each choice, and observing the resulting average revenues. For different bidder populations (varying bidder counts, risk sensitivity, and correlation of signals), different auction parameter values are found to maximize revenue.

Taking another tack, Phelps *et al.* explored the use of genetic programming to determine auction mechanism rules automatically.

Genetic programming (or GP), another form of evolutionary computation that is similar to GAS, evolves programs (or expressions) rather than the binary strings evolved in GAS. This makes automatic programming possible, and in theory allows even more flexibility and effectiveness in finding optimal solutions in the domain of concern. In GP, programs are traditionally encoded as *tree structures*. Every tree node has an operator function and every terminal node has an operand, making mathematical expressions easy to evolve and evaluate. With tree structures, crossover is applied on

an individual by simply switching one of its nodes with another node from another individual in the population. Mutation can replace a whole node in the selected individual, or it can replace just the information of that node. Replacing a node means replacing the whole branch. This adds greater effectiveness to the crossover and mutation operators [52].

Phelps *et al.* [38] demonstrated how GP can be used to find an optimal point in a space of pricing policies, where the notion of optimality is based on allocative efficiency and trader market power. In DA markets, there are two popular pricing policies: the k -DA pricing rule [43] and the uniform pricing policy. The former is clearly a discriminatory policy³⁹ and may be represented as:

$$p = k \cdot p_a + (1 - k) \cdot p_b$$

where $k \in [0, 1]$, and p_a and p_b are ask and bid prices. The latter executes all transactions at the same price, typically the middle point of the interval between the market ask and bid quotes. Searching in the space of arithmetic combinations of shout prices and market quotes including the above two rules as special cases, led to a complex expression that is virtually indistinguishable from the $k = 0.5$ version of the k -DA pricing rule. This shows that the middle-point transaction pricing rule not only reflects the traditional practice but also can be technically justified.

Noting that the performance of an auction mechanism always depends on the mix of traders participating in the mechanism, and both the auction mechanism and the trading strategies may adapt themselves simultaneously, Phelps *et al.* [36, ?] further investigated the use of co-evolution in optimizing auction mechanisms. They first co-evolved buyer and seller strategies and then together with auction mechanisms. The approach was able to produce outcomes with reasonable efficiency in both cases.

6.3 Evaluating market mechanisms

Phelps *et al.* proposed a novel way to evaluate and compare the performances of market mechanisms using heuristic strategy analysis [37].

Despite the fact that the performance of an auction mechanism may vary significantly when the mechanism engages different sets of trading agents, previous research on

³⁹That is transactions are cleared at different prices depending upon the prices of the matching bid and ask.

auctions analyzed the properties of DA markets using an arbitrary selection of homogeneous trading strategies. A more sound approach is to find the equilibria of the game between the participating trading strategies and measure the auction mechanism at those equilibrium points. As Sections 5.3 and 5.4.3 have discussed, the heuristic strategy analysis calculates equilibria among a representative collection of strategies. This makes the method ideal for measuring market mechanisms at those relatively stable equilibria.

The representative strategies selected by Phelps *et al.* included RE, PVT, and TT. The replicator dynamics analysis revealed that: (1) neither the CDA nor the CH mechanism is strategy-proof since TT is not dominant in either market; (2) increasing the number of agents in the CH led to the appearance of an equilibrium basin for an equilibrium near TT, which agreed with the conclusion drawn through the approximate analysis in [44] discussed in Section 3.2; and (3) the CH has higher efficiency than the CDA in the sense that the three equilibrium points⁴⁰ in the dynamics field for the CH all generate 100% efficiency while the only equilibrium⁴¹ for CDA produces 98% efficiency. One can interpret the small efficiency difference as justifying the NYSE's use of a CDA rather than a CH for faster transactions and higher volumes.

One avenue of future research is to combine this evaluation method with evolutionary computation to optimize DA mechanisms.

6.4 Adaptive auction mechanisms

Considering that the information about the population of traders is usually unknown to the auction mechanism, and many analytic methods depend on specific assumptions about traders, Pardoe and Stone advocated a self-adapting auction mechanism that adjusts auction parameters in response to past auction results [31].

Their framework includes an *evaluator* module, which can create an auction mechanism for online use, can monitor the performance of the mechanism, and can use the economic properties of the mechanism as feedback to guide the discovery of better parameter combinations. This process then creates better auction mechanisms that continue to interact with traders which are themselves possibly evolving at the same time. A classic algorithm for n -armed bandit problems, ϵ -greedy, is used in the evaluator module to make decisions on parameter value selection.

⁴⁰Each falls onto one of the three pure strategies, though the sizes of their basins vary.

⁴¹Pure RE strategy.

This work differs from previous work in the sense that here auction mechanisms are optimized during their operation while the mechanisms in the approaches discussed before find remain static and are assumed to perform well even when they face a set of traders that is different from those used in searching for the mechanisms.

6.5 Auction mechanism design competition

Following the TAC Classic and the TAC SCM competitions introduced in Section 5.5, a new competition called TAC CAT⁴² was run in the summer of 2007 in order to foster research on auction mechanism design. In TAC CAT, the software trading agents are created by the organizers of the competition, and entrants compete by defining rules for matching buyers and sellers and setting commission fees for providing this service. Entrants compete against each other in attracting buyers and sellers and making profits. This is achieved by having effective matching rules and setting appropriate fees that are a good trade-off between making profit and attracting traders.

I led the development of JCAT [30], based on Phelps's JASA,⁴³ to run as the game server. It provides various trading strategies, market selection strategies, and DA market mechanism frameworks to avoid entrants working from scratch. JCAT is also an ideal experimental platform for researchers to evaluate auction mechanisms in a competition setting.

⁴²CAT is not only the reverse of TAC, but also refers to *catallactics*, the science of exchanges.

⁴³JASA is a high-performance auction simulator that allows researchers in agent-based computational economics to run trading simulations using a number of different auction mechanisms. The software includes an implementation of the 4-heap algorithm in [55] and is designed to be highly extensible, so that new auction rules can easily be implemented. The software also provides base classes for implementing simple adaptive trading agents [33].

7 Summary

This report aims to provide an overview of the field of auction mechanism design and build the foundation for further research.

Auctions are markets with strict regulations where traders negotiate and make deals. An auction may be single-sided or double-sided depending upon whether only sellers or only buyers can make offers or whether both can. The four standard single-sided auctions—English auction, Dutch auction, first- and second-price sealed-bid auctions—have been the subject of traditional auction theory. Vickrey's pioneering work in this area led to the revenue equilibrium theorem that shows a seller can expect equal profits on average from all the standard types of auctions with a few assumptions about the bidders. Other researchers followed the approach and managed to extend the applicability of the theorem when the assumptions are relaxed.

Double-sided auctions, which are important in the business world, posed a bigger challenge due to the higher complexity of their structure and the interaction between traders. While classical mathematical approaches have continued to be successful in analyzing some simple types of double auctions, they have been unable to apply to more practical scenarios. Smith and others initiated experimental approaches and showed that double auctions, even with a handful of traders, may lead to high allocative efficiency and the transaction prices quickly converge to the expected equilibrium price. Subsequent experiments with human and/or artificial traders tried to explain what led to these desirable properties and tended to show that auction mechanisms played a major role, though the intelligence of traders had an effect as well.

Further work, on the one hand, introduced more and more complex trading strategies not only making higher individual profits but also improving the collective properties of auctions. On the other hand, different methods have been explored to design novel auction mechanisms. One approach is to evolve parameterized auction mechanisms based on evolutionary computation. Cliff *et al.* have found a new variant of continuous double auctions through evolving mechanisms that converge more quickly to equilibrium, and also exhibit higher efficiency than those previously known. Phelps *et al.* have explored the use of genetic programming and justified the traditional midpoint transaction pricing rule as optimizing efficiency while balancing trader market power. In addition to these off-line techniques for optimization through evolutionary computing, online approaches have been proposed to produce adaptive auction mechanisms, which, with dynamic trader populations, can continuously monitor and improve their performance.

With the understanding of this prior research work, what can be done further at the interface of computer science and economics include: obtaining more insights into double-sided auction mechanisms, inventing novel auction rules, and searching for optimal combinations of various kinds of policies, automatically producing desirable auction mechanisms.

References

- [1] Xuefeng Bao. Market making and the index futures market. *Perspectives*, 3(3), December 2001.
- [2] George W. Brown. Iterative solution of games by fictitious play. In T. C. Koopmans, editor, *Activity Analysis of Production and Allocation*, pages 374–376. Wiley, New York, 1951.
- [3] Andrew Byde. Applying evolutionary game theory to auction mechanism design. In *EC'03: Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 192–193, New York, NY, USA, 2003. ACM Press.
- [4] Kalyan Chatterjee and William Samuelson. Bargaining under incomplete information. *Operations Research*, 31(5):835–851, September–October 1983.
- [5] Dave Cliff. Evolution of market mechanism through a continuous space of auction-types. Technical report, Hewlett-Packard Research Laboratories, Bristol, England, 2001.
- [6] Dave Cliff. Evolutionary optimization of parameter sets for adaptive software-agent traders in continuous double auction markets. Technical report, Hewlett-Packard Research Laboratories, Bristol, England, 2001.
- [7] Dave Cliff and Janet Bruten. Minimal-intelligence agents for bargaining behaviours in market-based environments. Technical report, Hewlett-Packard Research Laboratories, Bristol, England, 1997.
- [8] Dave Cliff, Vibhu Walia, and Andrew Byde. Evolved hybrid auction mechanisms in non-ZIP trader marketplaces. Technical report, Hewlett-Packard Research Laboratories, Bristol, England, September 2002.
- [9] Vincent Conitzer and Tuomas Sandholm. Complexity results about Nash equilibria. Technical Report CMU-CS-02-135, School of Computer Science, Carnegie Mellon University, May 2002.
- [10] Vincent Conitzer and Tuomas Sandholm. Automated mechanism design: Complexity results stemming from the single-agent setting. In *The 5th International Conference on Electronic Commerce (ICEC'03)*, pages 17–24, Pittsburgh, PA, USA, September 2003.

- [11] Vincent Conitzer and Tuomas Sandholm. An algorithm for automatically designing deterministic mechanisms without payments. In *Third International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 1 (AAMAS'04)*, pages 128–135, New York City, NY, USA, July 2004.
- [12] Rajarshi Das, James E. Hanson, Jeffrey O. Kephart, and Gerald Tesauero. Agent-human interactions in the continuous double auction. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, Seattle, U.S.A., August 2001.
- [13] David Easley and John Ledyard. Theories of price formation and exchange in double oral auctions. In Friedman and Rust [15], chapter 3, pages 63–97.
- [14] Daniel Friedman. The double auction institution: A survey. In Friedman and Rust [15], chapter 1, pages 3–25.
- [15] Daniel Friedman and John Rust, editors. *The Double Auction Market: Institutions, Theories and Evidence*. Santa Fe Institute Studies in the Sciences of Complexity. Perseus Publishing, 1993.
- [16] Drew Fudenberg and David K. Levine. *The Theory of Learning in Games*. MIT Press, Cambridge, MA, 1998.
- [17] Steven Gjerstad and John Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22:1–29, 1998.
- [18] Dhananjay K. Gode and Shyam Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.
- [19] Dhananjay K. Gode and Shyam Sunder. Lower bounds for efficiency of surplus extraction in double auctions. In Friedman and Rust [15], chapter 7, pages 199–219.
- [20] Paul Klemperer. Auction Theory: A guide to the literature. *Journal of Economic Surveys*, 13(3):227–286, July 1999. <http://ideas.repec.org/a/bla/jecsur/v13y1999i3p227-86.html>.
- [21] Victor R. Lesser. Cooperative multiagent systems: A personal view of the state of the art. *IEEE Transactions on Knowledge and Data Engineering*, 11(1):133–142, 1999.
- [22] Ananth Madhavan. Trading mechanisms in securities markets. *Journal of Finance*, 47(2):607–641, 1992.

- [23] Eric Maskin and John Riley. Optimal auctions with risk averse buyers. *Econometrica*, 52(6):1473–1518, November 1984.
- [24] Randolph Preston McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56(2):434–450, 1992.
- [25] Paul R. Milgrom and Robert J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50:1089–1122, 1982.
- [26] Roger B. Myerson. Optimal auction design. *Mathematics of Operation Research*, 6:58–73, 1981.
- [27] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trade. *Journal of Economic Theory*, 29:265–281, 1983.
- [28] John F. Nash. Equilibrium points in n-person games. In *Proceedings of the National Academy of Sciences*, volume 36, pages 48–49, 1950.
- [29] James Nicolaisen, Valentin Petrov, and Leigh 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.
- [30] Jinzhong Niu. JCAT Tutorial — The Server Platform for TAC/CAT Competition. <http://jcat.sourceforge.net/>, 2008.
- [31] David Pardoe and Peter Stone. Developing adaptive auction mechanisms. *ACM SIGecom Exchanges*, 5(3):1–10, April 2005.
- [32] Simon Parsons, Mark Klein, and Juan Antonio Rodriguez-Aguilar. A bluffer’s guide to auctions. Technical report, Center for Coordination Science, Sloan School of Management, Massachusetts Institute of Technology, 2004. Research Note.
- [33] Steve Phelps. JASA — Java Auction Simulation API. <http://www.csc.liv.ac.uk/~sphelps/jasa/>, 2007.
- [34] Steve Phelps, Marek Marcinkiewicz, Simon Parsons, and Peter McBurney. Using population-based search and evolutionary game theory to acquire better-response strategies for the double-auction market. In *Proceedings of IJCAI-05 Workshop on Trading Agent Design and Analysis (TADA-05)*, 2005.

- [35] Steve Phelps, Marek Marcinkiewicz, Simon Parsons, and Peter McBurney. A novel method for automatic strategy acquisition in n-player non-zero-sum games. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'06)*, pages 705–712, New York, NY, USA, 2006. ACM Press.
- [36] Steve Phelps, Peter McBurney, Simon Parsons, and Elizabeth Sklar. Co-evolutionary auction mechanism design: a preliminary report. In *Proceedings of Workshop on Agent Mediated Electronic Commerce IV (AMEC IV)*, 2002.
- [37] Steve Phelps, Simon Parsons, and Peter McBurney. An evolutionary game-theoretic comparison of two double-auction market designs. In *Proceedings of Workshop on Agent Mediated Electronic Commerce VI (AMEC VI)*, 2004.
- [38] Steve Phelps, Simon Parsons, Elizabeth Sklar, and Peter McBurney. Using genetic programming to optimise pricing rules for a double auction market. In *Proceedings of the Workshop on Agents for Electronic Commerce,, Pittsburgh, PA, 2003*.
- [39] Chris Preist and Maarten van Tol. Adaptive agents in a persistent shout double auction. In *Proceedings of the 1st International Conference on Information and Computation Economies*, pages 11–18. ACM Press, 1998.
- [40] Alvin E. Roth and Ido Erev. Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior*, 8:164–212, 1995.
- [41] Michael H. Rothkopf and Sunju Park. An elementary introduction to auctions. Technical report, Rutgers Center for Operations Research, Rutgers University, January 2001.
- [42] John Rust, John H. Miller, and Richard G. Palmer. Behaviour of trading automata in a computerized double auction market. In Friedman and Rust [15], chapter 6, pages 155–199.
- [43] Mark Satterthwaite and Steven R. Williams. The Bayesian theory of the k -double auction. In Friedman and Rust [15], chapter 4, pages 99–123.
- [44] Mark A. Satterthwaite and Steven R. Williams. The rate of convergence to efficiency in the buyer's bid double auction as the market becomes large. *Review of Economic Studies*, 56(4):477–498, 1989.

- [45] Vernon L. Smith. An experimental study of competitive market behaviour. *Journal of Political Economy*, 70(2):111–137, April 1962.
- [46] Gerald Tesauro and Rajarshi Das. High-performance bidding agents for the continuous double auction. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*, pages 206–209, 2001.
- [47] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–27, 1961.
- [48] Vibhu Walia, Andrew Bye, and Dave Cliff. Evolving market design in zero-intelligence trader markets. Technical report, Hewlett-Packard Research Laboratories, Bristol, England, 2002.
- [49] William Walsh, Rajarshi Das, Gerald Tesauro, and Jeffrey O. Kephart. Analyzing complex strategic interactions in multi-agent systems. In Piotr Gmytrasiewicz and Simon Parsons, editors, *Proceedings of 2002 Workshop on Game-Theoretic and Decision-Theoretic Agents (GTDT-02)*, Edmonton, Alberta Canada, July 2002. AAAI.
- [50] William E. Walsh, David C. Parkes, and Rajarshi Das. Choosing samples to compute heuristic-strategy Nash equilibrium. In *AAMAS 2003 Workshop on Agent Mediated Electronic Commerce*, Melbourne, Australia, 2003.
- [51] Michael P. Wellman, Shih-Fen Cheng, Daniel M. Reeves, and Kevin M. Lochner. Trading agents competing: Performance, progress, and market effectiveness. *IEEE Intelligent Systems*, 18(6):48–53, 2003.
- [52] Wikipedia. Wikipedia. <http://wikipedia.org/>, 2008.
- [53] Robert Wilson. Equilibrium in bid-ask markets. In G. Feiwel, editor, *Arrow and the ascent of economic theory: Essays in Honor of Kenneth J. Arrow*, pages 375–414. McMillan, London, 1987.
- [54] Robert B. Wilson. Competitive bidding with disparate information. *Management Science*, 15(7):446–448, March 1969.
- [55] Peter R. Wurman, William E. Walsh, and Michael P. Wellman. Flexible double auctions for electronic commerce: Theory and implementation. *Decision Support Systems*, 1998.
- [56] Peter R. Wurman, Michael P. Wellman, and William E. Walsh. A parametrization of the auction design space. *Games and Economic Behavior*, 35:304–338, 2001.

[57] Peter R. Wurman, Michael P. Wellman, and William E. Walsh. Specifying rules for electronic auctions. *AI Magazine*, July 2002.