

Satterthwaite, Mark A.; Williams, Steven R.

Working Paper

The Double Auction Market: Institutions

Discussion Paper, No. 971

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Satterthwaite, Mark A.; Williams, Steven R. (1991) : The Double Auction Market: Institutions, Discussion Paper, No. 971, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/221330>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The Double Auction Market: Institutions,
Theories, and Experimental Evaluations

by

Mark A. Satterthwaite and Steven R. Williams*

November 13, 1991

1. Introduction

Supply, demand, and trade at a market clearing price are the most fundamental concepts of microeconomics. Marshall (1949) developed this analysis to describe properties of the outcome of trading that he believed to be essentially correct regardless of the actual institution through which trade is conducted. The supply-demand analysis is regarded in a different way in this paper. The "Marshallian cross" is used here as the institution for determining the market price-quantity pair, and the behavior of self-interested, imperfectly informed traders when confronted with this institution is analyzed. We study a call market, i.e., a market in which bids determine a demand curve, asks determine a supply curve, and all trades clear simultaneously at a market-clearing price.¹ This paper describes the institution, summarizes our theoretical knowledge of it, and discusses how those theoretical predictions may be confronted with experimental evidence.²

* Northwestern University. This material is based upon work supported by the National Science Foundation under Grant No. SES 9009546.

¹ Such an institution is commonly used to arrange trade in thin markets and for price discovery in more liquid markets at the opening of the trading day. See Schwartz (1988) for a discussion of call markets.

² While this paper mostly concerns our work and our joint work with Aldo Rustichini, we also wish to cite Chatterjee and Samuelson (1983), Wilson (1985), and Gresik and Satterthwaite (1989), upon which we built.

For simplicity, the theory is restricted to a trading environment consisting of m sellers and m buyers.³ Each seller has for sale a single, indivisible unit of a homogeneous good and each buyer is interested in purchasing one unit of the good. We assume that every trader has a reservation value for the good--cost c_i for a seller and value v_i for a buyer--that represents the value in money he places on a unit. Each trader privately knows his own reservation value. The assumption of private information is critical, for as Hayek (1945) emphasized, assuming instead that some "single mind" possesses all information relevant to trading trivializes the problem markets exist to solve.

The institution works as follows. At a set time each trader submits a sealed bid V_i if he is a buyer or offer C_i if he is a seller.⁴ The offers and bids (or asks) are arrayed to form "reported" supply and demand curves, a market clearing price p is selected, and units are exchanged among those sellers who offered less than p and those buyers who bid more than p . The market then disbands with no opportunity for recontracting.⁵ We call this institution a double auction (or DA) because both sides of the markets jointly determine the price through their offers/bids. In particular, for $k \in [0,1]$, the k-double auction (or k-DA) is the particular institution that selects $kb + (1-k)a$ as the price when $[a,b]$ is the interval in which a market-clearing price can be selected.

Figure 1 illustrates the institution for the case of $m = 3$, $k = 1$, and a specific realization of costs and values. What Marshall considered as the "true" supply and demand curves of the underlying economic environment are depicted by the step functions SS and DD : the three sellers had costs c_i equal to 0.30, 0.57, and 0.90 while the three buyers had values v_i equal to 0.73, 0.64, and 0.14. These curves depict willingness to trade as given by the traders' privately known reservation values. Sellers

³ The case in which there are different numbers of traders on each side of the market is considered in Williams (1991) and Rustichini, Satterthwaite and Williams (1990). Henceforth the latter reference will be referred to as RSW.

⁴ Because of the assumption of unitary supply/demand, there is no point to allowing traders in this restricted environment to submit multiple offers/bids or offers/bids together with a quantity. The institution--if not the theoretical results--described here is easily generalized to this richer form.

⁵ This is different from most real-world call markets in which some form of aftermarket typically exists. The possibility of trade in the aftermarket could influence behavior in the call market.

submitted offers C_i equal to their costs c_i . Buyers, however, attempted to manipulate the price by submitting bids less than their values; these bids of 0.53, 0.45, and 0.10 are represented on the figure by the dotted step function $D'D'$. Because $k = 1$ price is set at 0.53, which is the top of the interval of market clearing prices.

Three points should be noted about the example of Figure 1. First, relative to the reported supply and demand, the DA institution computes a market clearing price and assigns the available supply of goods to the m traders who reported the highest values for the units. In the example the price 0.53 is market clearing because the number of traders who submitted offers/bids of at least 0.53--thereby expressing that they placed a value of at least 0.53 on a unit--exactly equals the tendered supply of three units. Second, the first buyer with value 0.67 who bid 0.56 regrets ex post that he did not bid lower, for if he had bid 0.46 he would have still received a unit, but at a price of 0.46. This possibility of being able to influence price is what leads each buyer to bid less than his reservation value and causes the reported demand $D'D'$ to lie below true demand DD .⁶ Third, the tendency of buyers to choose their bids strategically may lead to inefficient outcomes. In the figure Pareto optimality requires two units to be traded: buyers one and two with values 0.73 and 0.64 should trade with sellers one and two with costs 0.30 and 0.57. Buyer two fails to trade because he attempted to manipulate price in his favor by bidding 0.45. As illustrated by the crossing of the true supply and demand curves SS and DD at quantity two, the outcome of trading would have been efficient if buyers had not acted strategically by reporting the bids that generated the demand curve $D'D'$.

Figure 1 reveals the difficulty of the problem that a trading institution must solve, even in such a simple setting. Because reservation values are private the institution must elicit traders' values from the traders themselves and then use those elicited values to allocate the units available for trade. But because the institution must use traders' reported values, traders may have an incentive to make offers/bids that are self-serving rather than honest. Thus the institution must strike a balance between

⁶ Sellers do not ask for more than their true value in this example because $k = 1$ and they can not influence the price at which they trade. This is the special case of the 1-DA that is discussed below. If $k < 1$, then sellers too would have an incentive to misrepresent their willingnesses to pay.

providing incentives to reveal private information accurately and fully using the information that is revealed to make the allocation.⁷

Theoretical research suggests that the k-DA makes this tradeoff remarkably well. This research, which we review in Sections 2, 3, and 4, has focused on "equilibrium" misrepresentation. It shows that when the number of traders is small (m equal to one or two) both somewhat efficient and very inefficient equilibria exist. No convincing theory has been developed that suggests which of these equilibria are most likely to occur in practice. As the number of traders grows, however, the set of equilibria shrinks, reported supply and demand quickly converge to true supply and demand, and the outcome of trading converges to the competitive price and efficient quantity. Our research thus supports Marshall's argument (1949, ch. II, 1-2) that true supply and demand essentially determines the outcome of trade even though traders do not have "thorough knowledge of the circumstances of the market" and thus try to obtain more favorable prices.

The theoretical analysis of the k-DA is radically complicated if even minor modifications of its rules are made so that it no longer uses a market-clearing price to mediate trade. Section 5 describes such a modification, which Kagel and Vogt (1991) used during the pilot phase of an experimental investigation of the k-DA, and discusses the analytical problems it causes. Finally, in Section 6 we review the testable propositions that the theory implies and discuss the possibility of using experimental methods to test these implications.

⁷ The k-DA takes the information traders reveal at face value and fully uses it in making its allocation of the available supply. In the mechanism design approach, the institution is carefully crafted for the specific trading environment so that traders are honestly reveal their private information. This is achieved, however, by not fully using the information that is revealed, e.g., two traders may reveal that they should trade, but the mechanism will not have them trade because to do so would destroy the incentives for truthful revelation. This characteristic of the mechanism design approach is illustrated by the optimal trading mechanisms that are constructed in Myerson and Satterthwaite (1983) and Gresik and Satterthwaite (1989).

2. Environment, Trader's Decision Problem, and the Bayesian Model of Equilibrium

Environment. There are m sellers and m buyers where $m \geq 2$, all of whom are risk neutral.⁸ If seller i succeeds in selling his indivisible unit of the trading good, then his utility is $p - c_i$ where p is the price he receives and c_i is the cost that he privately places on the unit. Otherwise his utility is zero. Analogously each buyer's utility is zero when he fails to trade and $v_i - p$ when he does, where v_i is the value he privately places on obtaining a unit.

Institution. Each trader submits a sealed offer (C_i) or bid (V_i) that is conditioned on his private knowledge of his cost c_i or value v_i . The offers and bids are arrayed in ascending order $s_{(1)} \leq s_{(2)} \leq \dots \leq s_{(2m)}$; $s_{(k)}$ is thus the k th order statistic of the $2m$ bids and offers. Price is set within the interval $[s_{(m)}, s_{(m+1)}]$ of possible market clearing prices at $p = (1-k)s_{(m)} + ks_{(m+1)}$ where k is a parameter selected from $[0, 1]$ that is fixed prior to the market being opened. Buyers who bid at least $s_{(m+1)}$ and sellers who offer at most $s_{(m)}$ trade. The only exception is if $s_{(m)} = s_{(m+1)}$ and the interval of market clearing prices is degenerate. In this event the quantity supplied may fail to equal the quantity demanded, and if necessary a fair lottery is run to determine who will trade among those traders who bid or offered p and who are on the long side of the market.

Of special note are the cases of $k = 1$ (as in Figure 1) and $k = 0$. Because price equals $s_{(m+1)}$ in the 1-DA, a seller cannot influence in his favor the price at which he actually trades. It is therefore in his best interest to submit his true reservation value as his offer, and reported supply is the same as true supply (Satterthwaite and Williams, 1989). Similarly, a buyer in the 0-DA has the incentive to submit his true value as his bid. For $k \in (0, 1)$, a trader on either side of the market can influence price in his favor. He therefore has an incentive to shade his offer/bid away from his true reservation value.

⁸ RSW includes the case in which all traders on the same side of the market have the same, possibly risk averse, utility function. Risk aversion of this form does not substantially complicate the analysis.

The trader's decision problem. Buyers and sellers' decision problems are symmetric.

Therefore consider a specific buyer, buyer one, let $v \equiv v_1$ be the value he places on the good, and let λ be the bid he is thinking about submitting. Buyer one's decision is risky, for the λ he submits affects both the probability that he will succeed in trading and, if he does trade, the expected price he will pay. Suppose he tests a value of λ by calculating the change in his expected utility if he raises his bid by a small amount $\Delta\lambda$. This calculation requires that he have a notion in the form of two probabilities concerning how the other m sellers and $m-1$ buyers are likely to bid.

Let $\zeta_{(1)} \leq \zeta_{(2)} \leq \dots \leq \zeta_{(2m-1)}$ be the random array of the bids and offers of the traders other than buyer one, noting in particular that $\zeta_{(k)}$ is the k th order statistic from the restricted sample that excludes buyer one's bid. This contrasts with $s_{(k)}$, which is the k th order statistic of the full sample of size $2m$ that includes his bid. To calculate the change in his expected utility buyer one needs to know, first, the probability that increasing his bid by $\Delta\lambda$ will cause his bid to jump over $\zeta_{(m)}$, for if he does jump over $\zeta_{(m)}$ he goes from not trading to trading and picks up a trade of value approximately $v-\lambda$. Second, he needs to know the probability that his bid λ is bracketed by $\zeta_{(m)}$ and $\zeta_{(m+1)}$, for if so, then increasing his bid by $\Delta\lambda$ will cause the price at which he trades to increase by $k\Delta\lambda$ to $(1-k)\zeta_{(m)} + k(\lambda+\Delta\lambda)$. His change in expected utility is therefore

$$(v-\lambda) Pr\{\zeta_{(m)} \in (\lambda, \lambda+\Delta\lambda)\} - k\Delta\lambda Pr\{\lambda \in (\zeta_{(m)}, \zeta_{(m+1)})\}. \quad (2.1)$$

where the first term is buyer one's expected gain from switching from being an unsuccessful bidder to being a successful bidder, and the second term is his expected loss from causing the price to rise by $k\Delta\lambda$.

Buyer one may in practice estimate these probabilities from the empirical distribution of other traders' bids and offers, and thence compute his optimal bid using (2.1). We return to this possibility below in Section 6 when we discuss experiments concerning the k -DA. A second possibility, which is more fruitful theoretically, is that buyer one deduces these probabilities from his knowledge of the

strategies of the other traders and his Bayesian prior beliefs concerning the distributions of sellers' costs and other buyers' values.

The Bayesian game model. Some definitions are needed for this second approach. A strategy is a function that maps a trader's reservation value into the offer/bid that he submits. Let $S(\cdot)$ and $B(\cdot)$ be the strategies that buyer one believes are being used by all sellers and all other buyers, respectively. Thus if seller i has cost c_i , he bids $C_i = S(c_i)$. Since traders' reservation values are private, buyer one is uncertain about other traders' costs and values. Let distribution F with density f represent buyer one's subjective beliefs concerning any seller's cost c_i and let distribution G with density g represent his subjective beliefs concerning any other buyer's value v_i . Assume buyer one regards the cost or value of any other trader as distributed independently of his own value v and the values and costs of every other trader. These assumption imply that the following three probabilities are well defined:

$K(\lambda)$ is the probability that if $m-1$ buyers bid using strategy B and $m-1$ sellers offer using strategy S , then exactly $m-1$ bids/offers are less than λ .

$L(\lambda)$ is the probability that if $m-2$ buyers bid using strategy B and m sellers offer using strategy S , then exactly $m-1$ bids/offers are less than λ .

$M(\lambda)$ is the probability that if $m-1$ buyers bid using strategy B and m sellers offer using strategy S , then exactly m bids/offers are less than λ . Note that $M(\lambda) \equiv \Pr\{\lambda \in (\zeta_{(m)}, \zeta_{(m+1)})\}$, as used in eq. (2.1).

Formulas for these probabilities can be found in Rustichini, Satterthwaite and Williams (1990, henceforth referred to as RSW).

Given S , B , F , and G , expression (2.1) implies that the formula for the marginal expected utility of buyer one with value v and bid λ is:

$$\frac{\partial U_B(v, \lambda)}{\partial \lambda} = (v - \lambda)h_B(\lambda) - kM(\lambda) \quad (2.2)$$

where $h_B(\lambda)$ is the density of $\zeta_{(m)}$. If $c = c(\lambda) \equiv S^{-1}(\lambda)$, $\dot{c} = c'(\lambda) = 1/S'[c(\lambda)]$, $v = v(\lambda) \equiv B^{-1}(\lambda)$, and $\dot{v} = v'(\lambda) = 1/B'[v(\lambda)]$ are all well-defined, then

$$h_B(\lambda) = mK(\lambda)f(c)c + (m-1)L(\lambda)g(v)v. \quad (2.3)$$

The first term on the right-hand side, when multiplied by $\Delta\lambda$, is the probability that buyer one's bid, if he increases it by $\Delta\lambda$, will jump over the bid of one of the m sellers, whose bid happens to be $\zeta_{(m)}$, while the second term is the marginal probability of passing one of the other $m-1$ buyers.⁹

Buyer one's best response strategy to (S,B) specifies for each of his possible reservation values v a bid λ that maximizes his expected utility conditional on v . With suitable regularity assumptions, if sellers are using strategy S and other buyers using strategy B , then buyer one can select his optimal bid by setting (2.2) equal to zero and solving for λ .

Bayesian Nash equilibrium. Following Harsanyi (1967-68), the equilibrium concept that has been used to study the k -DA theoretically is symmetric Bayesian Nash equilibrium. Suppose, exactly as described above for buyer one, the subjective beliefs of every trader concerning the reservation values of other sellers and buyers are described by the distributions F and G , and that this is common knowledge among all the traders. Suppose further that traders' reservation values are elements of $[0,1]$, that F and G are C^1 functions on $[0,1]$, and that the densities f and g are strictly positive on $[0,1]$. Consider a pair of strategies (S,B) . Together they are a symmetric Bayesian Nash equilibrium for the k -DA if (i) for each seller i strategy S is the best response strategy to the other $m-1$ sellers playing S and the m sellers playing B and (ii) for each buyer i strategy B is the best response strategy to m sellers playing S and the other $m-1$ buyers playing B . Asymmetric equilibria in which each trader plays a distinct strategy, B_i or S_i , may exist, but as of yet have proven intractable to analysis. "Equilibrium" in this paper thus means "symmetric equilibrium."

⁹ Pick a particular seller i . The probability i 's offer C_i is in $(\lambda, \lambda + \Delta\lambda)$ is $f(c)c \times \Delta\lambda$ because c_i is a random variable with density f , $C_i = S(c_i)$, and C_i has density $f(c)c \equiv f[C(C_i)]C'(C_i)$. Given $C_i \in (\lambda, \lambda + \Delta\lambda)$, $K(\lambda)$ is the probability it is $\zeta_{(m)}$. Therefore $K(\lambda)f(c)c\Delta\lambda$ is the probability (i) one's bid jumps over C_i and (ii) C_i is $\zeta_{(m)}$.

The reason for using the Bayesian Nash equilibrium concept in studying the k-DA institution is that it models rational equilibrium behavior in a setting with incomplete information. Information is incomplete in that each trader's reservation value is private and other traders only have beliefs about his value. Behavior is rational in that expected utility is maximized conditional on one's information.

There are, however, at least two drawbacks of the Bayesian Nash equilibrium concept. First, no well developed theory exists that explains how traders starting *de novo* jointly learn a set of equilibrium strategies (especially if more than one equilibrium exists). In other words, the theory does not include a process that results in the Bayesian Nash equilibrium. This poses a problem for the experimentalist, who receives little guidance from the theory concerning how equilibrium behavior is to be elicited from subject traders. Second, though the Bayesian Nash solution concept was created to model rational behavior in a setting with incomplete information, it makes the strong informational assumption that F and G are common knowledge among the traders in order to support the rationality of their behavior.

Elementary geometry of equilibrium strategies. Let $\underline{C} = \lim_{c \downarrow 0^+} S(c)$ and $\bar{V} = \lim_{v \uparrow 1^-} B(v)$. RSW show that for a given equilibrium (S,B), there exists numbers \bar{c} and \underline{v} such that: (i) a seller with value above \bar{c} or a buyer with value below \underline{v} trades with probability zero, while a seller with value below \bar{c} or a buyer with value above \underline{v} trades with positive probability; (ii) S is increasing over $[0, \bar{c}]$ and B is increasing over $[\underline{v}, 1]$; (iii) $\lim_{v \downarrow \underline{v}^+} B(v) = \underline{v} = \underline{C}$ and $\lim_{c \uparrow \bar{c}^-} S(c) = \bar{c} = \bar{V}$. This geometric relationship is depicted in Figure 2 for the case of a pair of continuous equilibrium strategies. In words, (i) implies that a buyer with too low a value or a seller with too high a cost will never trade. A trader with such a value feels no pressure to bid reasonably, e.g., a seller with cost c above \bar{c} may choose an arbitrarily large number as his offer. The intervals $[0, \bar{c}]$ and $[\underline{v}, 1]$ are thus called the intervals over which serious offers/bids are made.¹⁰

¹⁰ On a more technical note, B must be differentiable a.e. in $[\underline{v}, 1]$ and S must be differentiable a.e. in $[0, \bar{c}]$ because they are increasing over these intervals. This substantiates the first order approach outlined above for the case of arbitrary equilibria (S,B). The convergence result in the next section rests upon these first order conditions.

3. Convergence of Equilibria to Ex Post Classical Efficiency as the Number of Traders Increases

The introduction discussed the incentives of traders to try to manipulate the price in their favor by overbidding in the case of sellers and underbidding in the case of buyers. This strategic misrepresentation causes inefficiency because trades that should occur do not necessarily occur. The economist's intuition is that as the market's size (measured here by m) becomes large traders become essentially price-takers and the problems stemming from strategic misrepresentation vanish. Theorem 1 shows that this intuition is correct within the k -DA model. In particular, for every Bayesian Nash equilibrium, misrepresentation as measured by $S(c)-c$ and $v-B(v)$ is $O(1/m)$.

Theorem 1 (RSW, Th. 5.1). Consider any equilibrium (S,B) in which trade occurs with positive probability, every seller always offers at least as much as his cost, and every buyer bids at most his value. A number κ exists, whose value is a function of F and G but not of m or (S,B) , such that, for all $v \in (\underline{v},1]$ and $c \in [0,\bar{c})$,

$$S(c) - c \leq \frac{\kappa}{m} \tag{3.1}$$

and

$$v - B(v) \leq \frac{\kappa}{m}. \tag{3.2}$$

Additionally, $\underline{v} \leq \kappa/m$ and $\bar{c} \geq 1 - \kappa/m$.

Inequalities (3.1-2) bound misrepresentation for a serious offer or bid; the last sentence bounds the intervals over which misrepresentation cannot be bounded. Together these bounds describe the rate at which misrepresentation vanishes as the number of traders on each side of the market increases. It is

worth reiterating that this theorem applies to all Bayesian Nash equilibria of the k-DA, including ones in which S or B jump upward discontinuously at particular values of c or v.

Some intuition for this rate of convergence is obtained by outlining a proof of (3.2) in the case of a differentiable strategy B. Recall the buyer's first order condition (2.2). The first term in formula (2.3) for $h_B(\lambda)$ is the probability that the buyer by increasing his bid will jump over $\zeta_{(m)}$ and that $\zeta_{(m)}$ is a seller's offer. Omitting this nonnegative term from (2.2) produces the inequality

$$(v - \lambda) \leq \frac{kM(\lambda)}{(m-1)L(\lambda)\dot{v}} \quad (3.3)$$

where $B(v) = \lambda$. Now imagine the graph of the increasing function $\lambda = B(v)$ in the v, λ plane over its domain $[0,1]$. It lies below the $\lambda = v$ diagonal, reflecting underbidding by a buyer. The amount of misrepresentation, $v - B(v)$, is the vertical distance between the graph and the diagonal. Misrepresentation increases as v increases if and only if $B'(v) < 1$, or equivalently, if $\dot{v} > 1$. Suppose $v - B(v)$ is maximized at v' . Then at v' necessarily $\dot{v} \geq 1$, for otherwise a $v'' > v'$ would exist at which more misrepresentation occurs. Thus

$$v' - B(v') \leq \frac{kM(\lambda')}{(m-1)g(v')L(\lambda')\dot{v}} \leq \frac{k}{(m-1)g(v')} \frac{M(\lambda')}{L(\lambda')} \quad (3.4)$$

where $\lambda' = B(v')$ and \dot{v} is also evaluated at v' . Recall the definitions:

$L(\lambda)$ is the probability that if $m-2$ buyers bid using strategy B and m sellers offer using strategy S, then exactly $m-1$ bids/offers are less than λ .

$M(\lambda)$ is the probability that if $m-1$ buyers bid using strategy B and m sellers offer using strategy S, then exactly m bids/offers are less than λ .

Unless m is quite small, these two probabilities are essentially indistinguishable and approximately equal. This suggests that the ratio $M(\lambda)/L(\lambda)$ is bounded.¹¹ This, together with the assumption that $g(v)$ is positive on $[0,1]$, implies that $v' - B'(v') \leq \kappa/m$ for some $\kappa > 0$, as Theorem 1 states.

¹¹ Proving this is the main work of a formal proof of the theorem.

If each trader in a k-DA offered/bid his true reservation value (contrary to his self-interest), then the resulting allocation would be ex post classical efficient: no gains from trade could remain because a buyer would fail to trade if and only if his benefit v_i were less than the cost c_i of every seller who fails to trade.¹² Theorem 1's result that as m increases the amount of misrepresentation decreases as κ/m therefore implies that in expectation the k-DA approaches--but does not reach--ex post classical efficiency rapidly as the number of traders increases. Theorem 2 makes this implication precise.

For given m, k, F, G , and equilibrium (S,B), and for given realizations of the traders' reservation values, the gains from trade realized by a k-DA is

$$\sum_{i \in T_B} v_i - \sum_{i \in T_S} c_i \quad (3.5)$$

where T_B and T_S are the sets of buyers and sellers respectively who successfully trade. The expected gains from trade is the expected value of the realized gains from trade when traders' reservation values are distributed according to F and G . The potential expected gains from trade is the expected gains from trade if each trader were to honestly report his reservation value rather than following his equilibrium strategy. The relative efficiency of an equilibrium (S,B) is its expected gains from trade divided by the potential expected gains from trade.

Theorem 2 (RSW, Th. 6.1). Consider any equilibrium (S,B) in which trade occurs with positive probability, every seller always offers at least as much as his cost, and every buyer bids at most his value. A constant ξ exists, whose value is a function of F and G but not of m or (S,B), such that the relative efficiency of (S,B) is at least

¹² See Holmstrom and Myerson (1983) for a taxonomy of standards of efficiency under different informational assumptions.

$$1 - \frac{\xi}{m^2}.$$

An upper bound on the relative efficiency of any equilibrium (S,B) can be computed by constructing an optimal trading mechanism and calculating that mechanism's relative efficiency. Gresik and Satterthwaite (1989) contains an algorithm for constructing optimal mechanisms.

Theorem 2 is complementary to an important theorem of Wilson (1985): for sufficiently large m , equilibria of the k-DA are incentive efficient, provided that the equilibrium strategies (S,B) are differentiable and have bounded derivatives. An equilibrium of the k-DA is an incentive efficient trading mechanism if "it is false that it is common knowledge that another rule would improve some agents' expected gains from trade without reducing other's expected gains" (Wilson, 1985, pp. 1101). Wilson's theorem thus establishes that for large markets the k-DA's allocations are optimal within the constraints that the traders' private information imposes. Our Theorem 2 establishes that these constrained optimal allocations are in fact close to the classical, full information, optimal allocations.

Figure 3 provides a simple, supply-demand intuition for this result. In the figure the supply and demand curves drawn in solid lines represent buyers' and sellers' true market supply and demand with market clearing quantity $q' = 35$. Suppose for the sake of this example that these curves have slope 1 and -1 respectively. The dotted demand curves represent their reported aggregate supply and demand curves: sellers have over-reported the cost of their units by $\kappa/m = 5$ and buyers have under-reported the benefit of their units by an equal amount. If we assume that both supply curves have slope 1 and both demand curves have slope -1, then the market clearing quantity for the reported supply and demand curves is $q'' = q' - \kappa/m = 30$. The gains from trade that are unrealized as result of trading q'' units instead of q' units is the area of the shaded triangle; consistent with Theorem 2 its area is

$(\kappa/m)^2 = 25$. This is only an analogy, of course, for Figure 3 reflects neither the discreteness nor the incomplete information of our model.¹³

4. Computation and Multiplicity of Equilibria

The traders' first-order conditions define a system of ordinary differential equations that can not be solved in closed form, but are easy to solve numerically. Computation of equilibria is important for three reasons. First, computation suggests that a continuum of smooth equilibria exist when k is in $(0,1)$, but at most one smooth equilibrium exists when k takes the extreme values of either 0 or 1. This is important because we do not have general proof of existence of equilibria.¹⁴ Second, numerical examples suggest that a k -DA may be almost fully efficient with small values of m (e.g., $m = 6$ in the case of uniform F and G). Examples of such computations are presented in RSW and in Satterthwaite and Williams (1989). The bounds in Theorems 1 and 2 only suggest rapid convergence to efficiency; computation of κ and ξ in the Theorems provides bounds on misrepresentation and inefficiency that prove to very coarse in comparison to the actual performance of computed equilibria. Third, numerical computation produces explicit predictions of bidding behavior in the k -DA, which is surely helpful for experimental testing of this institution.

System of Equations Determining $(\dot{c}, \dot{\lambda}, \dot{v})$. Pick a point (c, λ, v) that satisfies the inequalities $0 < c < \lambda < v < 1$. Suppose (perhaps counterfactually) that equilibrium strategies (S, B) exist such that $S(c)$

¹³ It may be possible to deepen this analogy between our convergence to efficiency result and the area of the Harberger triangle by following Bulow and Roberts (1989) who showed that results of auction theory concerning a single seller and potential bidders with unknown reservation values have direct parallels in the theory of a discriminating monopolist in standard price theory. Their analogy replaces a bidder whose value is drawn from a distribution G on a continuum with a continuum of buyers whose values are distributed according to G .

¹⁴ For generic (F, G) Williams (1991) proved existence in the 1-DA of a piecewise smooth equilibrium.

$= \lambda$ and $B(v) = \lambda$. Pick a representative buyer i . Given that (S,B) is an equilibrium, then the buyer's first order condition $\partial U_B(v,\lambda)/\partial \lambda = 0$ from eq. (2.2) and (2.3) is satisfied at (c,λ,v) :

$$mK(\lambda)f(c) \dot{c} + (m-1)L(\lambda)g(v) \dot{v} - \frac{kM(\lambda)}{(v-\lambda)} \dot{\lambda} = 0 \quad (4.1)$$

or, in matrix form,

$$\begin{bmatrix} D_{BS}(c,\lambda,v) & D_{BB}(c,\lambda,v) & 0 \end{bmatrix} \times \begin{bmatrix} \dot{v} \\ \dot{c} \\ \dot{\lambda} \end{bmatrix} = \frac{kM(\lambda)}{v-\lambda} \dot{\lambda} \quad (4.2)$$

where $\dot{\lambda} = \partial \lambda / \partial \lambda = 1$. The first order condition of a representative seller is similar. Putting the two together with the tautology $\dot{\lambda} = 1$ gives a system of ordinary differential equations:¹⁵

$$\begin{bmatrix} D_{BS}(c,\lambda,v) & D_{BB}(c,\lambda,v) & 0 \\ D_{SS}(c,\lambda,v) & D_{SB}(c,\lambda,v) & 0 \\ 0 & 0 & 1 \end{bmatrix} \times \begin{bmatrix} \dot{v} \\ \dot{c} \\ \dot{\lambda} \end{bmatrix} = \begin{bmatrix} kM(\lambda)/(v-\lambda) \\ (1-k)N(\lambda)/(\lambda-c) \\ 1 \end{bmatrix} \dot{\lambda} \quad (4.3)$$

A smooth equilibrium (S,B) defines a solution curve to this ordinary differential equation by the formula $(c,\lambda,v) = (S^{-1}(\lambda), \lambda, B^{-1}(\lambda))$.

Computing Equilibria. The system (4.3) may be numerically integrated to obtain a solution using any of a number of standard techniques.¹⁶ A simple approach to computing solutions that represent an equilibrium is as follows. Pick an initial point $P_0 = (c_0, \lambda_0, v_0)$. Solve (4.3) at P_0 for the vector of derivatives $P_1' = (\dot{c}_1, \dot{\lambda}_1, \dot{v}_1)$ and pick a small, positive step $\Delta \lambda$. Compute a new point $P_1 = (c_0 + \dot{c}_1 \Delta \lambda, \lambda_0 + \dot{\lambda}_1 \Delta \lambda, v_0 + \dot{v}_1 \Delta \lambda)$. At P_1 solve (4.3) to obtain P_2' and then to compute P_2 . Continue this iterative process in the positive direction generating points P_2, P_3, P_4, \dots as long as neither event E_+ nor event e_+ occurs:

¹⁵ Formulas for D_{SS} , D_{SB} , and N can be found in RSW.

¹⁶ See Press et. al. (1986, ch. 15) for a sampling of effective numerical algorithms.

Event E_+ occurs at point $P_{n'} = (c_{n'}, \lambda_{n'}, v_{n'})$ if the inequalities $0 < c < \lambda < v < 1$ are violated.

Event e_+ occurs at point $P_{n'}$ if the vector of derivatives $P'_{n'+1} = (\dot{c}_{n'+1}, \dot{\lambda}_{n'+1}, \dot{v}_{n'+1})$ violates either $\dot{c}_{n'+1} > 0$ or $\dot{v}_{n'+1} > 0$.

If event e_+ occurs at point $P_{n'}$, then, because equilibrium strategies must be increasing, no smooth equilibrium goes through the initial point P_0 . Select a new initial point at which to restart the algorithm. If event E_+ occurs, return to P_0 , reverse the sign of $\Delta\lambda$ so as to generate points in the opposite direction, and iteratively generate points $P_{-1}, P_{-2}, P_{-3}, \dots$ as long as neither event E_- nor event e_- occurs. Events E_- and e_- are defined exactly as E_+ and e_+ except in the definition of e_- the vector of derivatives P'_{n-1} is tested against the inequalities.

If the process terminates at point $P_{n''}$ with event E_- occurring, then the points $\{(c_{n''}, \lambda_{n''}), \dots, (c_0, \lambda_0), \dots, (c_{n'}, \lambda_{n'})\}$ and $\{(v_{n''}, \lambda_{n''}), \dots, (v_0, \lambda_0), \dots, (v_{n'}, \lambda_{n'})\}$ numerically describe equilibrium strategies S and B respectively that go through the initial point P_0 , provided that at each of the points satisfaction of the first order conditions are sufficient for maximizing traders' expected utilities. A condition that guarantees that the first order conditions are sufficient for utility maximization is that F/f and $(G-1)/g$ are increasing functions on $[0,1]$ (Satterthwaite and Williams, 1989).

If event E_+ occurs at point $P_{n'} = (c_{n'}, \lambda_{n'}, v_{n'})$, then $c_{n'}$ approximates the upper endpoint \bar{c} of the interval over which a seller makes serious offers. Similarly, if E_- occurs at $P_{n''} = (c_{n''}, \lambda_{n''}, v_{n''})$, then $v_{n''}$ is approximately \underline{v} , i.e., the lower endpoint of the interval over which a buyer makes serious bids. Consequently E_- and E_+ are the correct tests at which to terminate the construction of (S,B) .

Figure 2 above graphs $S(c)$ and $B(v)$ for an equilibrium (S,B) for the case of $m = 2, k = 0.5$, and F and G are both the uniform distribution. The diagonal, $c_i = C_i$ and $v_i = V_i$ is drawn into the graph so the amount of misrepresentation that (S,B) involves can be judged by eye. If S and B involved no strategic misrepresentation, then the graphs of S and B would be congruent with the diagonal. Consequently the vertical distance between the diagonal and a point on the graph of S or B is $S(c_i) - c_i$ or $v_i - B(v_i)$ respectively.

Multiplicity and convergence of equilibria. For a given m , k , F , and G , the entire set of smooth equilibria can be represented by using the following computational strategy. Fix λ at the expected equilibrium price $\bar{\lambda}$ (0.5 when the F and G are the uniform distribution). Construct a grid initial starting points $(c, \bar{\lambda}, v)$ by selecting the cost c from the interval $(0, \bar{\lambda})$ and the value v from the interval $(\bar{\lambda}, 1)$. Calculate solutions to the system (4.3) for all the starting points and discard those solutions that are not equilibria. If the grid is made fine enough, then the resulting set of equilibria approximates the set of smooth equilibria.

Figure 4 shows the bundles of strategies S and B that result when this procedure is carried out for the case of $m = 2$, $k = 0.5$, and uniform F and G . Figure 5 repeats the procedure for $m = 4$. Comparison of the two figures makes obvious what Theorem 1 states must be the case: the maximal amount of misrepresentation of any of the strategies graphed for the $m = 4$ case is approximately half that in the $m = 2$ case.

Consider briefly the special case of the 1-DA. As noted earlier, a seller's dominant strategy in the 1-DA is honest reporting, i.e., $S^*(c) = c$. Here there is at most one smooth strategy B such that (S^*, B) is an equilibrium in the 1-DA (Williams, 1991). If F and G are the uniform distribution (for instance), then this equilibrium takes the simple form $S^*(c) = c$ and $B(v) = mv/(m+1)$. Uniqueness holds because, in addition to the first order conditions, there is an initial condition that B must satisfy-- $B(0) = 0$. A similar result holds for the 0-DA. A two dimensional bundle of equilibria exists when k is in the interval $(0, 1)$ because there is no well-defined initial condition for such k . Bayesian game theory thus provides a much more precise prediction for bidding behavior in the 0-DA and 1-DA than in the k -DA with $k \in (0, 1)$. Nevertheless, as Figures 4 and 5 show, the rapidity with which the equilibrium set shrinks as m increases essentially resolves, for large m , this indeterminacy.

5. A Modified Sealed Bid Double Auction

Kagel and Vogt (1991) introduced a subtly modified version of the sealed bid 1-DA. On initial analysis the changes seem innocuous, even helpful, for the purposes of experimental work. Further

analysis shows that this modified DA (MDA henceforth) has quite different properties than the standard 1-DA. Most notably, the price in the MDA may not clear the market; as a consequence sellers' incentives in the MDA are counterintuitive. We discuss the MDA here to highlight the virtue of the k-DA.

Rules of the MDA. The rules of the MDA are identical to the rules of the 1-DA except for one change in the algorithm for selecting the price. In the 1-DA, price is set at $s_{(m+1)}$ whether the trader who offered/bid $s_{(m+1)}$ is a seller or a buyer. The MDA sets price at $s_{(m+1)}$ if it is a buyer who bids $s_{(m+1)}$. If it was a seller who offered $s_{(m+1)}$, then price is set equal to the bid of the buyer whose bid exceeds $s_{(m+1)}$ by the smallest amount. In other words, the price is set equal to the smallest of the bids of those buyers who get to buy. Trade still occurs between buyers who bid at least $s_{(m+1)}$ and sellers whose offers were no more than $s_{(m)}$.

For example, suppose $m = 3$, sellers' offers are 0.43, 0.49, and 0.81, and buyer's bids are 0.64, 0.32, and 0.11. The MDA sets price at $s_{(5)} = 0.64$. This contrasts with the 1-DA, in which price would equal $s_{(4)} = 0.49$. Both mechanisms prescribe that the only trade is between seller one with $C_1 = 0.43$ and buyer three with $V_3 = 0.64$. The price 0.64 is not a market clearing price, for seller two who offered 0.49 is not allowed to trade at the price 0.64 even though his offer indicates that he would earn a positive return at that price.

Incentives to misrepresent. While a buyer in the 1-DA sets the price only if his bid equals $s_{(m+1)}$ in the set of all $2m$ bids and offers, his bid is the price in the MDA if it equals $s_{(m+j)}$ for $j \geq 1$ and no bids lie between $s_{(m)}$ and $s_{(m+j)}$. This increased likelihood of setting price increases the expected reward from underbidding, which means that a buyer in the MDA has an incentive to bid further below his value v_i than would be the case with the 1-DA.

The change in rules also affects sellers' incentives. The MDA eliminates truthful reporting as a dominant strategy and gives a seller an incentive to make an offer C_i below his cost c_i . As in the 1-DA, a seller has no incentive to overbid because he can not affect the price at which he trades. Sellers have an incentive to underbid, however, because the MDA's price is not market clearing. When a price

is chosen in the MDA that exceeds a value that clears the market (i.e., whenever $s_{(m+1)}$ is a seller's offer), there is an excess supply available at the price, with only those sellers whose offers were below $s_{(m+1)}$ being able to trade. This possibility of an excess supply provides a seller with an incentive to underbid so as to include himself among those sellers who get to sell. To see this, consider again the example above in which sellers' offers are 0.43, 0.49, and 0.81, buyers' bids are 0.11, 0.32, and 0.64, price is set at $s_{(5)} = 0.64$, and seller two who offered 0.49 is excluded from trade. After the bids/offers are opened seller two regrets that he did not underbid by making an offer less than 0.43, for if he had done so, then he would have traded profitably at the unchanged price of 0.64. Underbidding, of course, also includes the possibility of trading at a price below one's value; in equilibrium, a seller in the MDA weighs these two effects in determining the optimal amount by which he underbids.

Characterization and computation of equilibria. Characterization and computation of equilibria is straightforward for the k-DA; it is quite the opposite for the MDA, which implies that it is hard to obtain testable predictions. The source of the difficulty is that a seller (and similarly, a buyer) who checks that his offer λ is optimal must compute not only the likelihood that λ is less than $\zeta_{(m)}$, but also the likelihood that price will be set equal to $\zeta_{(m+1)}$, $\zeta_{(m+2)}$, ..., or $\zeta_{(2m-1)}$. This implies that each trader's first order condition requires global knowledge of other traders' strategies, not just their local properties around the offer/bid λ he is testing. This contrasts with the k-DA in which each trader's first order condition uses only local information about other traders' strategies (as shown by (4.3)). Consequently no simple characterization of equilibria through the first order conditions is possible and the computational technique we used for the k-DA of constructing traders' equilibrium strategies through a series of small steps fails for the MDA.

Our experience in attempting to compute equilibria for the particular case of uniform F and G has caused us to question whether pure strategy, symmetric equilibria even exist in the MDA.¹⁷

¹⁷ We were able to compute ϵ -equilibria for the MDA by constructing a sequence of buyers and sellers' strategies through myopic adjustment. Specifically, we began with S_0 and B_0 , which were not an equilibrium pair. We computed S_1 , the best response of a seller to other traders playing S_0 and B_0 . To construct S_1 we computed the seller's optimal offer for a variety of costs and then fitted a Chebyshev approximation to the resulting cost-bid pairs. Given S_1 , we computed B_1 , a buyer's best

Computationally we found that if $S(\cdot)$ has the property that, for all $c > 0$, $S(c)$ is strictly positive, then the best response $S^*(\cdot)$ of a seller to (S,B) is, for c sufficiently close to zero, to offer $S^*(c) = 0$. But if $S(\cdot)$ has the property that, for c sufficiently close to 0, $S(c) = 0$, then the best response $S^*(\cdot)$ to (S,B) is, for all c , to offer $S^*(c) > 0$. Hence no symmetric pure strategy equilibria appears to exist.¹⁸ We are uncertain if mixed strategy equilibria exist.

Comments. Despite the similarity of the MDA with the 1-DA, it appears to be substantially more difficult to analyze. It is not obvious that results concerning the 1-DA have analogues for the MDA. Our lack of theoretical or computational results cripples the MDA's usefulness as an institution to be tested experimentally, for there are no solid predictions concerning what behavior it should induce from subjects.

6. Experimental Testing of the Bayesian Nash DA Theory

From an experimental viewpoint the theory of the k -DA has at least three virtues. First, the theory is mainstream in that the Bayesian-Nash solution concept is (for better or worse) currently the most widely accepted method for modeling strategic behavior when information is incomplete. Second, the model on which the theory is based is easily translated into the lab for it fully specifies the generation of preferences, the information traders receive and the actions they may take, and the algorithm for computing price and making a final allocation. Third, the theory generates a number of testable hypotheses about buyer and seller behavior in symmetric equilibria and how it changes as m , the number of traders on each side of the market, increases. The three most obvious are (i) sellers offer no less than their reservation values and buyers bid no more, with offers strictly less when $k < 1$ and

response to other traders playing S_1 and B_0 . We continued this process until it converged to an ϵ -equilibrium, which in our very limited experience occurred quickly. We do not report these computational results here because the results of our computations indicate that no pure strategy equilibria exist for the MDA.

¹⁸ This nonexistence conjecture may have little or no relevance to experimental investigations of the MDA because ϵ -equilibria do appear to exist that fail to be best responses only for values of c close to zero. A seller with such a value of c is almost certain to sell his unit. Thus an experimental subject is unlikely to worry much about his offers for such values of c .

bids strictly more when $k > 0$; (ii) the maximal amount by which sellers overbid and buyers underbid relative to their reservation values is proportional to β/m where $\beta > 0$, and (iii) the relative efficiency of any equilibrium must be greater than $1 - \gamma/m^2$ where $\gamma > 0$.¹⁹ Moreover, given the distributions F and G , the constants β , γ , and γ' are calculable as is the entire set of equilibria with smooth strategies (S,B). If the theory is true in the positive sense, then the behavior of experimental subjects should conform to all three implications.

A difficulty in devising a credible test. Despite the cleanness of the model, the directness of its transfer into the laboratory, and the precision of the predictions, it is unclear how to falsify the theory. The problem is this. The Bayesian theory of double auctions posits sophisticated behavior on the part of traders within a game of incomplete information. How to behave in a k-DA is unlikely to be transparent to an inexperienced participant. He must make careful inferences from noisy data about other traders' strategies and then optimize his own bidding behavior against their (imperfectly understood) behavior. As a consequence an experimental subject may need to participate in a k-DA a very large number of times before he can accomplish the learning necessary to play in accordance with the theory's predictions.

During the early stages of an experimental subject's experience with the k-DA, he almost certainly does not understand what is his optimal behavior. He may therefore fall back on his life experience with economic institutions and adopt a behavior that, by analogy, seems appropriate. This initial, rule-of-thumb behavior does not falsify the theory because the possibility remains that the subject will with sufficient experience change his behavior in a way that brings it into conformance with the theory. Real traders, as opposed to experimental subjects, have enormous experience and therefore may consistently behave in agreement with the theory even though their learning was in fact painfully slow.

¹⁹ In addition, for any m and any (F,G) , an upper bound on the relative efficiency can be computed by constructing an optimal mechanism and calculating its relative efficiency. Gresik and Satterthwaite (1989) describe the construction of an optimal mechanism. Computations they present suggest that, for any (F,G) , a positive $\gamma' < \gamma$ exists such that the relative efficiency of the optimal mechanism (and thus of any equilibrium of the k-DA) can not exceed $1 - \gamma'/m^2$.

Consequently the question any experimental test of the theory must confront, particularly if it obtains negative results, is: How much experience must subjects have for the test to be credible?

As an example of this, Kagel and Vogt (1991) had subjects play the MDA that is described above in Section 5. Sellers have an incentive in the MDA to make offers that are less than their reservation values. They in fact made offers that were greater than their reservation values and thus acted in accordance with the standard intuition most individuals carry into trading situations in which price is negotiated: a seller should ask for more than what the object is worth to him.²⁰

Experimentalists understand that learning is critical, but have not made much progress in defining its magnitude or speed.²¹ Nevertheless at least three complementary approaches may be taken to improve the likelihood that experimental results contradicting k-DA theory will be regarded as credible. First, theory may be developed that gives a sense of how many repetitions a subject would have to play a particular k-DA situation in order to at least have a reasonable chance of learning equilibrium behavior. Second, the experiment may have to be designed to make the learning process for the subjects easier. Third, subjects can be matched against computer players in order to allow controlled measurements of how fast learning proceeds in the k-DA environment.

Theory-based approaches to improved credibility. To our knowledge even the most elementary notions of how long learning should take in the k-DA settings have not been developed. For example, suppose a seller has reservation value c and is trying to ascertain if changing his bid to λ'' from λ' would increase his expected utility. How many repetitions of the DA would he have to play in order to reject the hypothesis that λ' and λ'' yield identical expected utilities? Answers to questions such as this would give a some sense of how quickly a fully rational double auction participant could correct

²⁰ Kagel and Vogt also had subjects participate in the 1-DA. Though sellers have a dominant strategy in the 1-DA of truthful reporting, they also tended to follow the standard intuition of asking for more than their reservation values.

²¹ See Kagel's review (1991, Section 1.6) of learning within private value auctions. The evidence is weak and he comes to no firm conclusions concerning either its speed or its effectiveness. More definitely, within a series of public goods experiments Palfrey and Rosenthal (1991) report that subjects demonstrated a very limited ability to learn implicitly critical statistics describing other players' strategies.

nonequilibrium behavior and would place a lower bound on the number of repetitions an actual, boundedly rational, experimental subject would have to play the k-DA before he could be expected to play in accordance with equilibrium predictions.²²

An experimental design that aids learning. The simplest way to design a k-DA experiment for a market with m traders on each side of the market is to take $2m$ experimental subjects, assign half to be sellers and half to be buyers, and have them play n repetitions of the k-DA. Each repetition played is an independent event in that each time every trader independently draws a new reservation value from the appropriate underlying distribution F or G . In this setup each seller implicitly defines his entire strategy $S(\cdot)$ as he selects offers C_i in response to different values of his cost c_i . Simultaneously each buyer is implicitly defining his entire strategy $B(\cdot)$. This is clearly a difficult task for traders to do: discover an optimal S or B even as other traders are changing their strategies as they too search for an optimal strategy.

One way of making the learning task easier is to ask each experimental subject to define only one point of his strategy. Specifically, assemble a large group of $2M$ subjects, with $M \gg m$, assign half to be buyers and half to be sellers, and give each one a permanent reservation value so that the distribution of the M sellers' costs approximates F and distribution of M buyers' values approximates G . Conduct n repetitions of the k-DA by drawing independently for each repetition a random market of m buyers and m sellers from the two large pools of M sellers and M buyers. Within each repetition the individual trader faces a market of $2m-1$ other traders whose reservation values are distributed precisely as they would be if new reservation values were drawn for each repetition. Since each individual seller keeps the same cost c_i for all n repetitions of the DA, he is required only to choose a single point $S(c_i)$ on his strategy, a much easier task than computing the entire function $S(\cdot)$. Similarly, each buyer only has to choose the optimal bid for one value.

²² Such calculations may be particularly illuminating for the k-DA because in it trade occurs at a uniform price that is determined by at most two of the $2m$ bids and offers. Consequently a nonoptimal bid/offer only occasionally affects a trader's realized payoffs, i.e., the k-DA's feedback to traders is generally weak and sporadic.

Computer simulated players and learning best responses. A more tightly controlled approach can be taken to investigate how rapidly--if at all--a trader learns an optimal strategy. Suppose, for instance, a seller played the k-DA against a set of $2m-1$ computer simulated players who are programmed to play equilibrium strategies. Each repetition he would submit an offer simultaneously with the computer submitting the other $2m-1$ bids/offers that it calculates by substituting randomly generated reservation values into equilibrium strategies (S,B). The subject could be assigned a new, independently drawn cost c_i each repetition as in the standard design discussed above, or he could keep a single c_i across all repetitions as in the modified design.

The main advantage of having the subject play against a computer is that the subject faces a stationary problem, i.e., the subject does not face a distribution of bids and offers from the other $2m-1$ traders that is changing as they revise their strategies on the basis of their own learning.²³ Positive results from running this tightly controlled design would not establish that traders tend to play Bayesian equilibrium strategies in the k-DA, for this design does not address how a group of traders can simultaneously learn equilibrium behavior. Negative results, however, would be evidence against the theory's positive validity because the ability of a single trader to learn his best response against equilibrium behavior is certainly a necessary condition for a set of traders to learn equilibrium behavior.

7. Concluding Comment

The Bayesian Nash equilibrium theory of the k-DA provides some theoretical justification that the predictions of supply/demand analysis may be valid even in small markets with imperfectly informed, strategic traders. This is consistent with experimental research on a variety of other institutions for organizing trade. Unestablished, however, is whether or not this theory positively describes the behavior of experimental subjects in the k-DA. Testing the theory is difficult because optimal behavior in the k-DA is far from transparent to experimental subjects. Only if experienced

²³ A disadvantage of unknown importance is that a subject may play differently if he knows that he is playing against computer simulated traders as opposed to real traders.

subjects fail to behave in accordance with the theory's predictions will it be falsified. But no criterion exists to decide if a group of subjects is sufficiently experienced that data from their actions can falsify the theory. A convincing test of this theory thus seems to require a more detailed theory of learning in Bayesian games.

References

- Bulow, J., and J. Roberts. 1989. The simple economics of optimal auctions. J. of Political Econ. 97: 1060-1090.
- Chatterjee, K. and Samuelson, W. 1983. Bargaining Under Incomplete Information. Operations Research 31: 835-851.
- Gresik, T. and M. Satterthwaite. 1989. The rate at which a simple market converges to efficiency as the number of traders increases: An asymptotic result for optimal mechanisms. J. of Econ. Theory 48 (June): 304-332.
- Harsanyi, J. 1967-68. Games with incomplete information played by Bayesian players, Parts I, II, and III. Management Science 14: 159-182, 320-334, 486-502.
- Hayek, A. 1945. The use of knowledge in society. The American Economic Review 35, 519-530.
- Holmstrom, B. and R. Myerson. 1983. Efficient and durable decision rules with incomplete information. Econometrica 51: 1799-1819.
- Kagel, J. 1991. Auctions: A survey of experimental research. Xeroxed, March. To appear in the Handbook of Experimental Economics.
- Kagel, J. and W. Vogt. 1991. Buyer's bid double auctions: preliminary experimental results. mimeo., University of Pittsburgh.
- Marshall, A. (1949). Principles of Economics. London, The MacMillan Press.
- Myerson, R. and Satterthwaite, M. 1983. Efficient mechanisms for bilateral trade, Journal of Economic Theory 29: 265-281.

- Palfrey, T. and H. Rosenthal. 1991. Testing game-theoretic models of free riding: New evidence on probability bias and learning. In Laboratory Research in Political Economy, ed. T. Palfrey, pp. 239-268. Ann Arbor: University of Michigan Press.
- Press, W., B. Flannery, S. Teukolsky, and W. Vetterling. 1986. Numerical Recipes: The Art of Scientific Computing. Cambridge, England: Cambridge University Press, 1986.
- Rustichini, A., M. Satterthwaite, and S. Williams. 1990. Convergence to price-taking behavior in a simple market. D.P. 914, Center for Mathematical Studies in Economics and Management Science, December.
- Satterthwaite, M. and S. Williams. 1989. The rate of convergence to efficiency in the buyer's bid double auction as the market becomes large. Review of Economic Studies 56, 477-498.
- Schwartz, R. 1988. Equity Markets: Structure, Trading, and Performance. New York: Harper & Row, 1988.
- Williams, S. 1991. Existence and convergence of equilibria in the buyer's bid double auction. Rev. Economic Studies 58, 351-374.
- Wilson, R. 1985. Incentive efficiency of double auctions", Econometrica 53: 1101-1115.

Captions to Figures

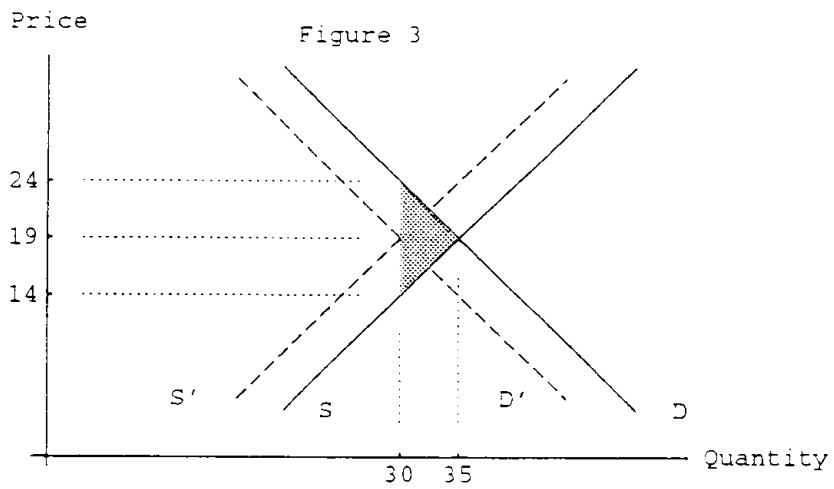


Figure 1. An example of a DA with $m = 3$ and $k = 1$. SS and DD are the true supply and demand curves and $D'D'$ is the demand curve strategically reported by buyers. Sellers report their true supply curve because $k = 1$. The price relative to reported supply and demand is 0.53; one unit is traded at that price.

Figure 2. Equilibrium strategy pair (S,B) for $m = 2$, $k = 0.5$, and (F,G) uniform on $[0,1]$. Note that $\underline{c} = \underline{v} = 0.172$ and $\bar{c} = \bar{v} = 0.828$.

Figure 3. True supply and demand are the solid lines S and D . Reported supply and demand are the dashed lines S' and D' . The shaded triangle represents the gains from trade lost because S' and D' are reported.


Figure 4. Equilibrium strategies (S,B) for $m = 2$, $k = 0.5$, and (F,G) uniform. The equilibrium (S,B) pairs here approximate the full range of equilibrium behavior for $m = 2$.

Figure 5. Equilibrium strategies (S,B) for $m = 4$, $k = 0.5$, and (F,G) uniform. The equilibrium (S,B) pairs here approximate the full range of equilibrium behavior for $m = 4$.



$S(c), B(v)$

Figure 4



$S(c), B(v)$

Figure 5

