Dhananjay K. Gode and Shyam Sunder
Graduate School of Industrial Administration, Carnegie-Mellon University, Schenley Park, Pittsburgh PA 15213

SEVEN
Lower Bounds for Efficiency of Surplus Extraction in Double Auctions

The double auction is a robust institution for efficient extraction of consumer and producer surplus in a variety of market environments. This conclusion is based on evidence gathered in a large number of laboratory markets populated by profit-motivated human traders, usually students (see Smith,[14] Plott,[6] and Friedman[4] in this volume for literature reviews). The efficiency of surplus extraction in double auctions may derive from the characteristics of this institution, from characteristics of trader behavior, or perhaps from interactions between the two. Virtually all economic experiments of the past have focused on the effect of varying economic institutions or environments on the performance of the market, holding profit-motivated behavior of traders unchanged. Since profit maximization is a maintained hypothesis in this literature, there has been an unchallenged inclination that it plays an important role in the tendency of double auctions to extract most of the consumer and producer surplus.

Gode and Sunder[5] substituted profit-motivated human traders by "budget-constrained zero-intelligence machine traders" in their double auction experiments. These traders are simple computer programs that generate random bids (or asks) subject to a no-loss constraint. The focus of that experiment was on examining the effect of controlled variation in trader behavior on the efficiency of the market. They found that imposition of a no-loss constraint (prohibiting traders from buying above their redemption values or selling below their cost) is sufficient to attain over 98% efficiency, even if traders, stripped of all rationality, submit random bids and asks. These results suggest that surplus extraction may largely be a property of the double auction institution, independent of the trader behavior.

This paper is an attempt to determine the lower bounds for expected surplus extraction efficiency of an idealized double auction populated with budget-constrained zero-intelligence (ZI) traders. The ultimate goal of this effort is to gain insights into the factors responsible for the high efficiency of some markets such as double auctions. The double auction appears to be too complex a game to yield a clear game-theoretic solution; its properties are easier to analyze with traders who act randomly.[1] We use the technique of using zero-intelligence traders to map the sensitivity of the lower bound of the expected efficiency to market parameters, continuously clearing procedures, relative number of extra-marginal traders, and rounds of bidding. The worst-case expected efficiency turns out to be 81% for "synchronized" double auctions and 75% for "continuously clearing" double auctions (these variations of double auction will be defined in the paper). The expected surplus extraction efficiency of continuous markets is lower because they permit a higher chance for extra-marginal traders to displace the intra-marginal traders. As the relative proportion of extra-marginal traders declines, expected efficiency of double auction converges to the close neighborhood of 100%. Some 50–100% of surplus is extracted in the first round of bidding in a synchronized double auction, declining sharply in subsequent rounds.

These analytical and simulation results confirm, and provide a better understanding of, the empirical and simulation results presented in Gode and Sunder.[5] A key insight is the crucial role of the tradeoff between the probability of an efficiency-reducing transaction and the magnitude of the resultant efficiency reduction in defining the environment for which the expected efficiency of double auction attains its lower bound. Second, the extra-marginal shapes of supply and demand functions affect the expected efficiency of double auction through their effect on both the probability as well as the magnitude of efficiency reduction, even though the theoretical equilibrium prediction is independent of the location of extra-marginal units.

The first section of the paper specifies a simple environment (i.e., demand and supply), trader behavior, and an idealized form of the institution labelled "synchronized" double auction. Expected efficiency of this auction with zero-intelligence traders is derived in Section 2. Section 3 examines the sensitivity of this result to the relative proportion of extra-marginal traders using computer simulations. Section 4 derives the efficiency results for a "continuously clearing" version of double auction which is a closer approximation of field and laboratory versions of this institution. Section 5 maps the time profile of surplus extraction in a synchronized double auction. Finally, Section 6 summarizes the results, and discusses its implications for the source of efficiency of double auctions.

[1]See Easley and Ledyard[2] for another alternative approach to study of double auctions. They specify boundedly rational, but non-strategic rules of trader behavior to examine the institutional properties.

THE MODEL

DEMAND AND SUPPLY

Each buyer has the right to buy up to one unit that has a given redemption value between 0 and 1. In the base case we consider first, there is one buyer with redemption value 1 and an infinite number of buyers with redemption value of β $(0 < \beta < 1)$. Similarly, each seller has the right to sell up to one unit with a given variable cost between 0 and 1. There is one seller with cost of 0 and an infinite number of sellers with cost $1 - \alpha$ $(0 < \alpha < 1)$. Figure 1 shows the resulting demand and supply functions, an equilibrium price range from β to $(1 - \alpha)$, and an equilibrium quantity of 1.

Two features of the demand and supply configuration we have chosen deserve comment. The presence of only one intra-marginal trader on each side of the market precludes the possibility of competition within the intra-marginal traders on each

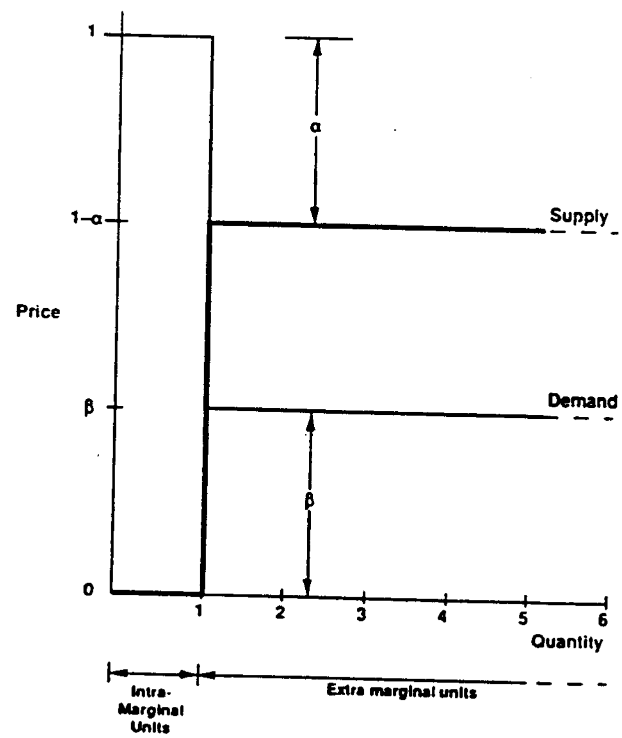

FIGURE 1 **Demand and Supply Functions.**

side. Competition amongst the intra-marginal buyers (sellers) only serves to exclude the extra-marginal buyers (sellers) from the market by raising the bids above (lowering the asks below) the redemption values (costs) of the extra-marginal buyers (sellers) more quickly. Thus competition among intra-marginal traders lowers the chances of displacement of intra-marginal traders by extra-marginal traders. Since such displacement is the only source of inefficiency in markets with ZI traders, we have used only one intra-marginal trader on each side to arrive at the lower bounds for efficiency. Efficiencies with intra-marginal competition would be higher.

The use of flat supply and demand functions in the extra-marginal region makes it easier to examine the effect of changing the costs (redemption values) of extra-marginal sellers (buyers) on efficiency. Presence of extra-marginal units at multiple redemption values or costs adds complexity without further insights. Finally, the use of a single intra-marginal trader and multiple extra-marginal traders at the same value or cost allows us to study the impact of changing the number of extra-marginal traders relative to the number of intra-marginal traders.

"ZERO INTELLIGENCE" TRADERS

Whenever a buyer has an opportunity to make a bid (to be defined by market rules discussed below), it generates random bids distributed independently and uniformly between 0 and the redemption value of its current unit. The imposition of an upper limit of redemption value on bids amounts to a budget constraint and prevents buyers from buying things they cannot afford to pay for. Similarly a seller generates (whenever it has an opportunity to do so) random offers distributed independently and uniformly between the cost of its current unit and 1. This ensures that the traders would not incur a loss. The lower limit of cost imposed on sellers' asks also has the effect of imposing a similar budget constraint on sellers. The support of the bids/asks generated by these traders does not change in response to the level of highest outstanding bid (current bid) or lowest outstanding ask (current ask).

MARKET RULES FOR SYNCHRONIZED DOUBLE AUCTION

In the base case we assume that the market operates as a ***synchronized double auction***.[2] All buyers and sellers are simultaneously solicited for bids and asks until each provides a bid or ask which does not violate its budget (i.e., the no-loss) constraint. The highest bid and the lowest ask are designated as the current bid and the current ask, respectively. If the current bid and the current ask cross, these units are traded in a binding transaction.[3] If they don't cross, solicitation from all

[2] See Rust, Miller, and Palmer,[10] in this volume, for an implementation of the synchronized double auction of which this is an idealized description.

[3] The price at which such transactions are booked can, depending on the market rules chosen, lie anywhere in the range between the crossed bid and ask. Since we are not concerned with the behavior of prices in this article, we leave this market rule unspecified.

traders is repeated. The bid/ask improvement rule is applied to calculate the new current bid/ask. This means that the current bid can be updated in a later round only by a higher subsequent bid and the current ask can be updated only by a lower subsequent ask. A transaction cancels all unaccepted bids and asks. This process is repeated until expiration of the prespecified time allowed for the period.

This particular idealization of double auction is labelled "synchronized" because in each round of bid/ask solicitation every trader's bid or ask is on the board before the highest bids and the lowest asks are allowed to close a transaction. In a later section of the paper, we analyze "continuously clearing" double auctions that bear greater resemblance to laboratory and field institutions.

EXPECTED EFFICIENCY

Expected efficiency of a synchronized double auction populated by budget-constrained zero-intelligence traders (with one intra-marginal buyer and seller each, and an infinity of extra-marginal buyers and sellers) is given by:

$$\begin{cases} 1-(a-\alpha-\beta)\left(\dfrac{\alpha^2}{1-\alpha(1-\alpha)}+\dfrac{\beta^2}{1-\beta(1-\beta)}\right) & \text{for} \quad \alpha+\beta\leq 1\,, \\ 1 & \text{for} \quad \alpha+\beta>1\,. \end{cases} \tag{1}$$

(See Appendix for Proof.)

The upper panel of Figure 2 shows expected efficiency as a function of the redemption values of extra-marginal buyers (β) and costs of extra-marginal sellers ($1-\alpha$). In the lower panel, the vertical scale has been expanded to show clearly the shape of the surface. Efficiency is 100% if $\alpha=\beta=0$ or if $\alpha+\beta\geq 1$. When $\alpha=\beta=0$, the extra-marginal buyers' costs are equal to the values of intra-marginal sellers, and the value of extra-marginal buyers' costs are equal to the values of intra-marginal sellers, making it impossible for the extra-marginal traders to enter the market. Given a sufficient number of rounds to submit bids and asks, intra-marginal traders transact their respective units to yield 100% efficiency.

If the demand and supply do not intersect in this market (i.e., $\alpha+\beta\geq 1$), all surplus is necessarily extracted in this market. Efficiency-reducing transactions can take place only when the cost of extra-marginal sellers exceeds the redemption value of extra-marginal buyers (i.e., $\alpha+\beta<1$).

The minimum possible value of expected efficiency in α,β-plane is 80.84%, attained at two points—($\alpha=0,\beta=0.639$) and ($\alpha=0.639,\beta=0$). The minimum is driven by two considerations. Usually, extra-marginal units cannot displace intra-marginal units because the budget constraint prevents these ZI traders from submitting bids that are high enough (or asks that are low enough) to transact. However, variability of transaction prices means that some transactions do take place at prices which are accessible to the extra-marginal traders. The extra-marginal units closer to the equilibrium price have a greater chance of displacing the intra-marginal units because bids are bounded above by the demand function (and asks are bounded below by the supply function) due to imposition of the budget constraint. When such displacement does occur, the units closer to the equilibrium price cause only a small loss of surplus extracted. Therefore, the *expected* loss of surplus due to displacement of intra-marginal units by extra-marginal units close to equilibrium is relatively small. On the other hand, extra-marginal units far away from the equilibrium price can have a big impact on surplus extracted whenever they are able to displace intra-marginal units in trading. But the chances of such displacement become increasingly remote as the distance of such units from the equilibrium price increases. Again, these units, too, have little effect on expected efficiency of the double auction. The maximum reduction in expected efficiency derives from intermediate units with a moderate effect on the magnitude of surplus extraction and only a moderate chance of displacing the intra-marginal units.

This minimum expected efficiency of 80.84% understates the expected efficiency one may expect to observe in a synchronized double auction on average. If we assume that parameters α and β for a particular auction are realizations of random variables drawn from uniform and independent distributions between 0 and 1, every point of the surface shown in Figure 1 would be equally likely. The average height of this surface and, therefore, the average expected efficiency over α, β-plane is 95.8%. This efficiency is only a few percentage points below the efficiencies observed in double auctions populated with profit-motivated human traders. This result seems to support Gode and Sunder's[5] conjecture that extraction of virtually all the surplus is a characteristic of the double auction, independent of trader motivation, intelligence, or learning. If allocative efficiency is to be equated with smartness, then smartness must be attributed to this institution itself; traders need not be smart to attain high efficiency in double auctions.

SENSITIVITY TO THE NUMBER OF EXTRA-MARGINAL TRADERS

Expression (1) and Figure 2 have been derived for the extreme case of a single intra-marginal trader and an infinity of extra-marginal traders on either side of the market. Since we do not have analytical expressions for finite values of N (the number of extra-marginal traders on either side of the market), the firm dark line for $N \rightarrow \infty$ in Figure 3 has been computed from Expression (1) and represents

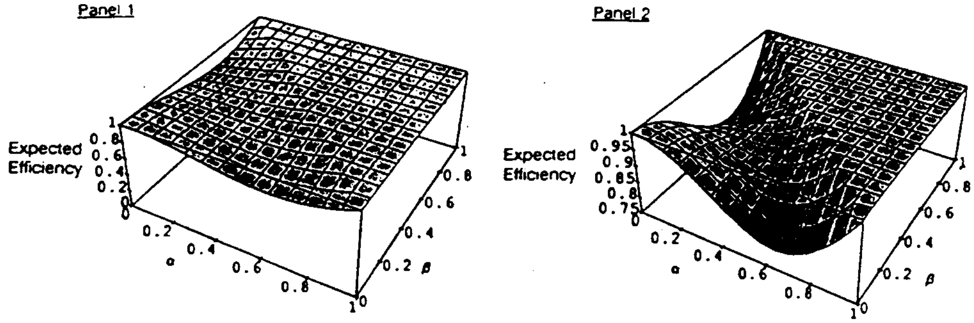

FIGURE 2 Expected efficiency of a synchronized double auction with zero-intelligence traders. Expected efficiency $= 1-(1-\alpha-\beta)(\alpha^2/(1-\alpha(1-\alpha))+\beta^2/(1-\beta(1-\beta)))$. Cost of one intra-marginal unit $= 0$. Cost of infinite extra-marginal units $= 1-\alpha$. Value of one intra-marginal unit $= 1$. Value of infinite extra-marginal units $= \beta$.

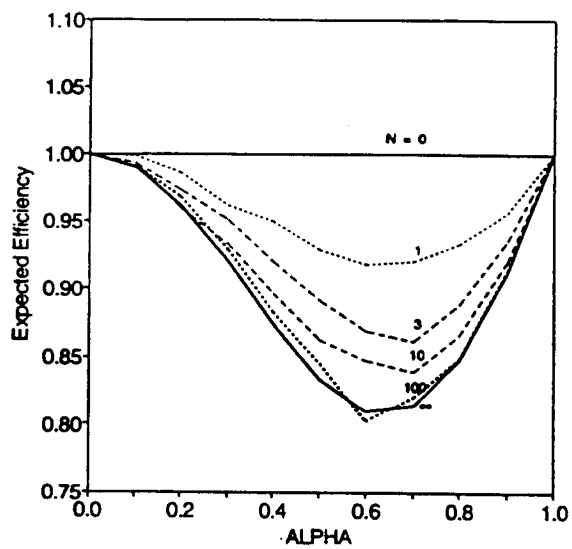

FIGURE 3 Effect of the number of extra-marginal traders on expected efficiency of synchronized double auctions (simulations with zero-intelligence traders). Value of extra-marginal buyers (β) is kept fixed at zero throughout, while the cost of extra-marginal sellers $(1-\alpha)$ is varied between 0 and 1. Expected efficiencies for $N = 0$ and $N \to \infty$ have been computed; all others have been simulated.

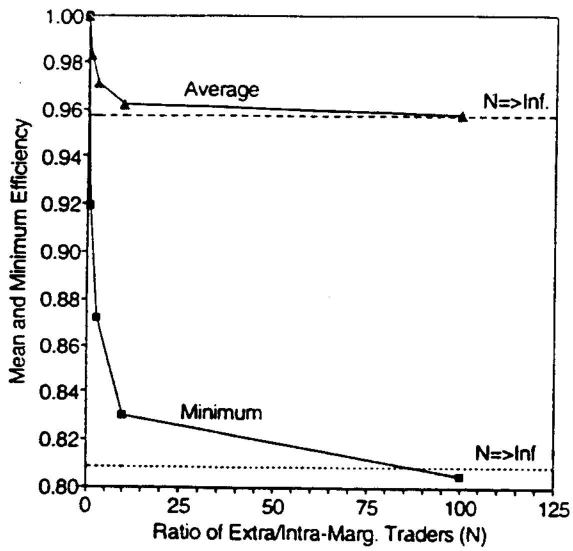

FIGURE 4 Effect of number of extra-marginal traders on minimum and mean efficiency over α, β-plane. Assuming $\alpha(= 1-$ cost of extra-marginal sellers) and β (value to extra-marginal buyers) are distributed uniformly and independently between 0 and 1. Simulations with zero-intelligence traders. Efficiencies for $N = 0$ and $N \rightarrow \infty$ have been computed; all others have been simulated.

the intersection of the efficiency surface with $\beta = 0$ plane in Figure 2. For finite values of N, we don't have closed form expressions; the broken lines in Figure 3 represent average efficiencies computed from 1,000 iterations of synchronized double auctions populated by budget-constrained zero-intelligence traders for $\beta = 0$ and $\alpha = 0, 0.1, 0.2, \ldots, 1$. When there are no extra-marginal traders, it is not possible for an efficiency-reducing transaction to take place. Since intra-marginal units necessarily get traded, efficiency of a market with $N = 0$ is necessarily 100%. As the number of extra-marginal trader increases without bound, expected efficiency converges to the lower bound specified by Expression (1) and shown in Figure 2. With fewer extra-marginal traders, the chances of efficiency-reducing transactions that may involve such traders also decrease, raising the expected efficiency of the market.

Simulated synchronized double auctions show that the general shape of the expected efficiency surface with respect to α and β given in Figure 2 for $N \rightarrow \infty$ remains unchanged for smaller values of N. Minimum expected efficiency is attained when either α or β is zero, and the other parameter is between 0.6 and 0.7. The minimum level attained drops sharply when N is increased from 0 to 1; further increases in N bring further reductions of diminishing magnitudes in the minimum level of expected efficiency attained. Virtually all the drop has been attained by the time N reaches 100. This behavior of minimum expected efficiency is confirmed in the lower curve in Figure 4.[4]

Figure 4 shows the expected efficiencies of synchronized double auction averaged over the α, β-plane for various values of N. The asymptotic value of this average for $N \to \infty$ is computed by integrating Expression (1) and is shown by a dashed line. Averages for finite values of N are obtained from simulations over an α, β-grid of fineness 0.1 and are shown in solid triangles. Average expected efficiency drops at a decreasing rate from 1 (for $N = 0$) to 0.958 (for $N = 100$) with the lower bound of 0.957 as $N \to \infty$. The total loss of this average expected efficiency is no more than 4.3% in the worst possible case. The dotted line and solid squares plot the corresponding minimum expected efficiency levels attained in the α, β-plane as a function of N.

Since there is only one intra-marginal buyer and seller each in these markets, N can be interpreted as the *ratio* of extra- to intra-marginal traders.[5] For most experimental markets, this ratio is rarely greater than two or three, corresponding to mean expected efficiency (over the α, β-plane) of 97% or higher. In naturally occurring markets, larger values of this ratio could be observed. Even then, the mean expected efficiency over the parameter space cannot drop much below 96%.

CONTINUOUSLY CLEARING DOUBLE AUCTIONS

Most double auctions in the field and laboratory differ from the synchronized double auction examined above in an important respect: the timing of their bids/asks is determined by the free will of individual traders, and a transaction is completed as soon as a bid and an ask cross, without waiting for the remaining traders in the market to submit their bids/asks. How sensitive are the findings of the preceding section to this variation in the rules of a double auction market?

There are several possible ways of formally modeling a continuously clearing double auction. A precise specification of the market rules is necessary for unambiguous analysis. We use a simple specification: all traders are randomly sampled (without replacement) to submit a bid or ask. Sampling without replacement in this case means that every trader gets a chance to submit one bid/ask before any trader is able to submit a second bid/ask, and so on. A bid (ask) becomes the current or market bid (ask) if it is higher (lower) than the standing market bid (ask). As soon as the inside quotes match or cross, a transaction is executed without waiting for

[4]Note that at one point in Figure 3 ($\alpha = 0.6, \beta = 0$), the expected efficiency for $N = 100$ is less than the lower bound of expected efficiency with these parameters when $N \to \infty$. This is possible because the lower bound is a computed value from Expression (1) while the expected efficiency for $N = 100$ is a sample mean from 1000 iterations, with a positive sampling error.

[5]At this point, it is not clear if the ratio of extra- to intra-marginal traders is an appropriate parameterization of the problem. We intend to conduct further studies to find out.

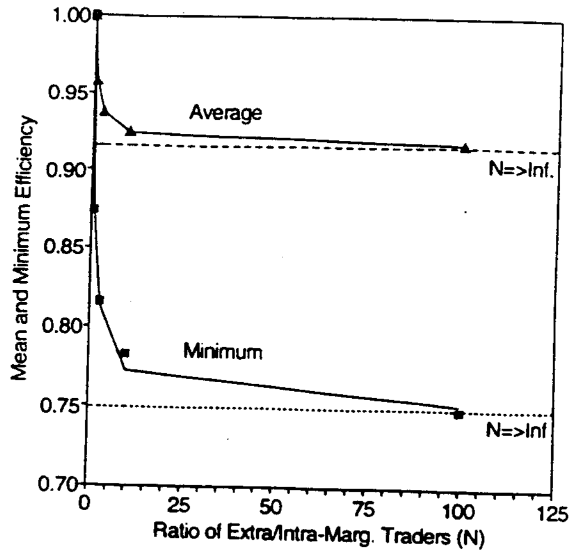

FIGURE 5 Effect of number of extra-marginal traders on minimum and mean efficiency over α, β-plane. Assuming α (= 1− cost of extra-marginal sellers) and β (value to extra-marginal buyers) are distributed uniformly and independently between 0 and 1. Simulations of continuous DA with zero-intelligence traders. Efficiencies for $N = 0$ and $N \to \infty$ have been computed; all others have been simulated.

the remaining buyers(sellers) to submit their bids (asks). If a transaction does not take place in the first round of bidding, further rounds are continued until it does. As soon as a transaction is executed, all unaccepted bids and asks are cancelled (i.e., the market bid is set to zero and the market ask is set to its maximum permissible value). In this model, market outcomes depend only on the order in which a given set of bids and asks are received; the time distribution of arrivals does not matter. Since the simulations presented in this section involve identical machine traders, the order in which each trader submits its bid/ask is randomized within each round by the above procedure.

Figure 5 shows the sensitivity of minimum and mean expected efficiency of continuous auctions to the value of N from 1,000 simulations. The solid squares plot the minimum observed value of mean expected efficiency over the α, β-plane as a function of N. The solid triangles plot the observed value of mean expected efficiency averaged over the α, β-plane from 1000 iterations of continuous double auctions.

These simulation results are closely approximated by the following expression:

$$1 - \frac{N}{1+N}(\alpha + \beta)(1 - \alpha - \beta) \tag{2}$$

where N is the number of extra-marginal buyers *and* the number of extra-marginal sellers. In order to understand this approximation, it is helpful to consider the four terms of this expression in turn. If there were no surplus-reducing transactions (i.e., transactions involving the extra-marginal traders), the expected efficiency would be 1, the first term of the expression. Whenever one extra-marginal buyer buys from the intra-marginal seller (and a surplus of β is realized), the intra-marginal buyer is necessarily forced to trade with one extra-marginal seller (to realize a surplus of α). Thus a surplus-reducing transaction results in a lost surplus of $(1 - \alpha - \beta)$, the last term of the expression. $N/(N+1)$ is the probability that an intra-marginal seller will face an extra-marginal buyer before it faces an intra-marginal buyer, and β is the approximate probability that the ask submitted by these zero-intelligence traders will cross each other.[6] Similarly, $N/(N+1)$ is the probability that the intra-marginal buyer will face an extra-marginal seller before facing an intra-marginal seller, and α *is* the approximate probability that they will consummate a transaction. Subtracting the expected loss of surplus from 1 yields expression (2) for approximate expected efficiency of the continuously clearing double auction. The shape of this approximate expected efficiency surface as a function of α and β as $N \to \infty$ is shown in Figure 6. The minimum value and mean value (over the α, β-plane) of this approximate expected efficiency as functions of N are shown in Figure 5 by the respective firm lines. The approximation is sufficiently precise that most of the simulated points (dark squares and triangles) appear, in spite of their sampling errors, to have been joined together by these lines.

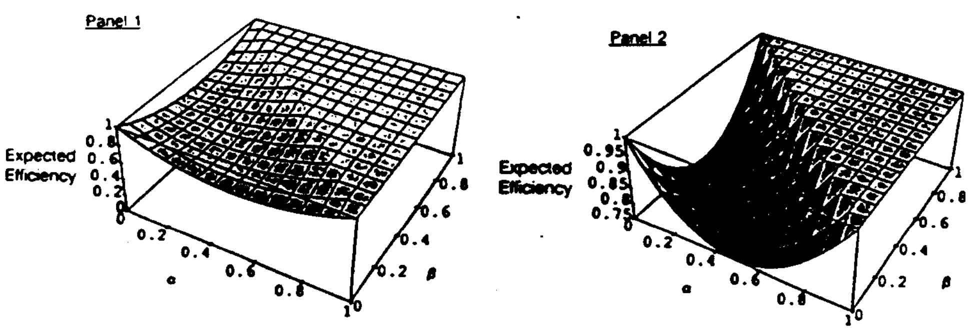

FIGURE 6 Effect of the number of continuous double auctions with zero-intelligence traders. Expected efficiency $= 1 - (N/(N+1)(\alpha + \beta)(1 - \alpha - \beta)$. Cost of one intra-marginal unit $= 0$. Cost of infinite extra-marginal units $= 1 - \alpha$. Value of one intra-marginal unit $= 1$. Value of infinite extra-marginal units $= \beta$.

[6]The reason this probability is approximate is that it ignores the consequences of no trade in the first round of bidding.

A comparison of Figure 6 with Figure 2 reveals some similarities and differences. The surface in Figure 6 is below the surface in Figure 2 everywhere. This means that the expected efficiency of synchronized double auction dominates the expected efficiency of continuous double auction at every point and on average. The maximum efficiency of 100% is retained for $\alpha = \beta = 0$ and for $\alpha + \beta \geq 1$. The minimum expected efficiency of continuous double auction is reached at all points of a line defined by $\alpha + \beta = 0.5$ (instead of just two points in Figure 2). The minimum level drops from 0.808 in Figure 2 to 0.75 in Figure 6. In interpreting this finding, we should be careful about several factors. First, these differences may narrow or disappear in markets populated by profit-motivated intelligent traders. Second, continuous double auctions may have the advantage of faster price discovery in dynamically changing markets and our present analysis does not include consideration of such factors. While the New York Stock Exchange could be thought of as a continuously clearing market, some of its new challengers are being designed as call markets somewhat similar to the synchronized double auction. Many stock markets around the world, especially those with less liquidity, operate as call markets. In a dynamic environment with new information, the price discovery role of markets may be important. It is possible that the higher static surplus extraction efficiency of the synchronized double auctions documented in this paper may be traded off against the possibly higher dynamic informational efficiency of continuous double auctions. For markets with sufficient liquidity, informational advantages of the continuous auctions may become sufficient to overcome their static surplus extraction disadvantages. The results obtained here point to these and other conjectures and directions of further investigation.

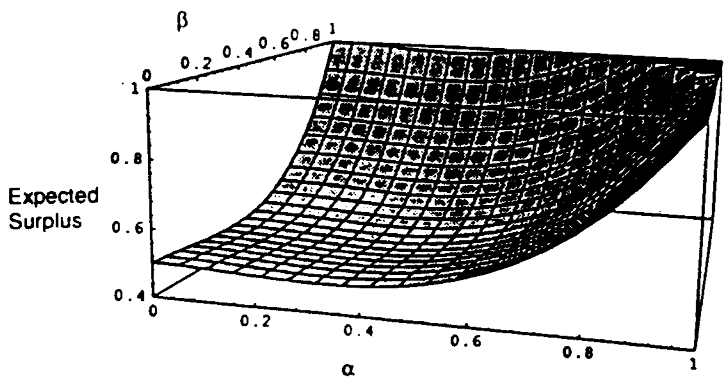

FIGURE 7 Expected surplus extraction in the first round of bidding in synchronized DA with zero-intelligence traders. Expected surplus $= 0.5 + \alpha^2(\alpha - 0.5) + \beta^2(\beta - 0.5)$. Cost of one intra-marginal unit $= 0$. Cost of infinite extra-marginal units $= 1 - \alpha$. Value of one intra-marginal unit $= 1$. Value of infinite extra-marginal units $= \beta$.

TIME PROFILE OF SURPLUS EXTRACTION

For a synchronized double auction, probabilities given in Table 1 (see the Appendix) can be used to derive the amount of surplus expected to be extracted in the first round of bidding when the number of extra-marginal traders $N \to \infty$:

$$0.5 + \alpha^2(\alpha - 0,5) + \beta^2(\beta - 0.5) \text{ for } \alpha + \beta \leq 1. \quad (3)$$

The shape of this function is shown in Figure 7 (only the part of the surface in the lower left side, $\alpha + \beta < 1$, is valid). Surplus extracted in the first round attains its maximum of 1 at the two corner points ($\alpha = 1, \beta = 0$) and ($\alpha = 0, \beta = 1$). The first-round expected efficiency is at its minimum of 0.463 at ($\alpha = 1/3, \beta = 1/3$) which is a little over half of the total expected efficiency of 0.905 at this point given

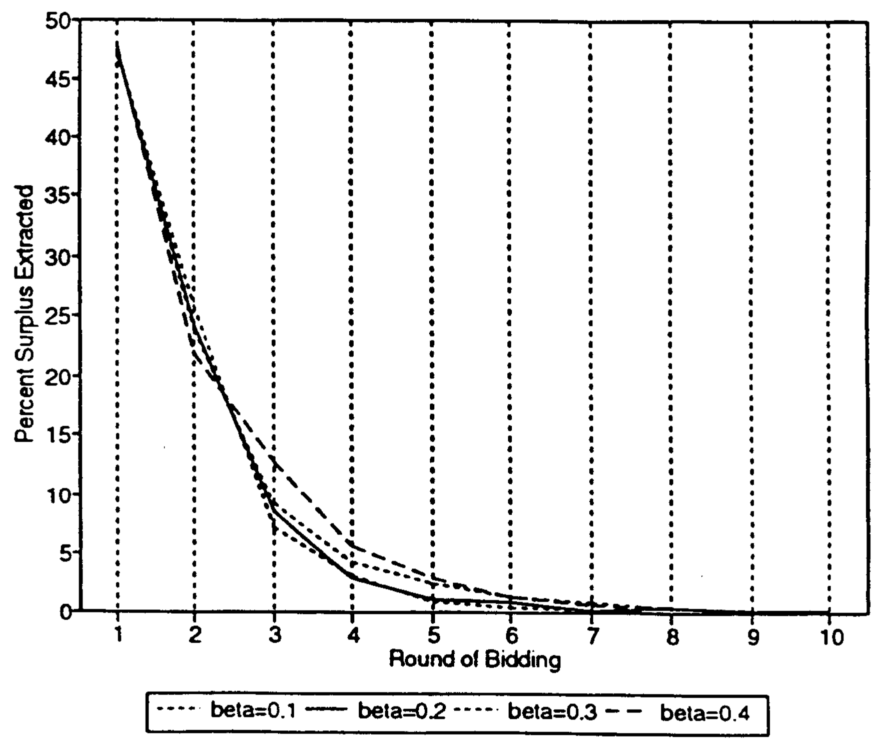

FIGURE 8 Time profile of surplus extraction in synchronized double auctions (by round).

by Expression (1). Thus, anywhere from 50–100% of the surplus extracted in a synchronized double auction is extracted within the first round itself.[7]

While we do not have the expressions for the amount of surplus extracted in the subsequent rounds, it must decline at a rapid rate because such a high proportion is extracted in the first round itself. Figure 8 shows the time profile of surplus extraction in the first ten rounds of such auctions for $\alpha = 0.5, \beta = 0.1, 0.2, 0.3$, and 0.4 and $N \to \infty$ from 1000 iterations of a computer simulation.[8] As expected, the surplus extraction declines sharply within a few rounds, little surplus being extracted after the first four or five rounds of bidding. Repeated bidding is an essential feature of double auctions. The first round efficiency of the double auction could be compared to the efficiency of similar market institutions that are constrained to a single round of bidding.

CONCLUDING REMARKS

The assumption of utility-maximizing traders in economic theory is often criticized because it ignores the well-documented cognitive limitations of human beings. However, the implications of this discrepancy between facts and assumptions about human behavior for the predictions of economic theory have remained controversial. Zero-intelligence (ZI) traders serve as a lower benchmark of intelligence. The results obtained here demonstrate the robustness of certain theoretical predictions when the assumption of individual rationality is relaxed in the extreme.

The design of exchanges populated by the identical, high-speed ZI traders requires a precise specification of market rules. This poses interesting problems for modeling and implementation of these markets. In the computerized markets of the laboratory, the speed of the underlying hardware and software exceeds the speed of human response by several orders of magnitude, giving rise to a wide range of micro-level issues in modeling of trader behavior and implementation of market institutions. We do not address these important micro-level design issues.

The demand and supply configuration used for the analysis can serve as a useful framework for studying the impact of extra-marginal traders on market efficiency. We have used a simple model of double auction, populated by zero-intelligence traders, to arrive at the following conclusions. Whether, or to what extent, these conclusions will hold in more complex double auction settings is an open issue.

[7] Recall that we are considering an environment in which there is only one intra-marginal unit. Consequently, the first round efficiency of synchronized double auction given by expression (3) is the same as the efficiency of buyer's bid double auction (see Satterthwaite and Williams,[11] in this volume). For double auctions with multiple intra-marginal units, the meaning of the "first round" would have to be appropriately redefined.

[8] In these simulations of synchronized double auctions it is surprisingly easy to let $N \to \infty$ simply by setting the highest extra-marginal bid to β and the lowest extra-marginal ask to $1 - \alpha$ in each round of bidding.

1. The expected efficiency of double auctions has a lower bound. Even when they are populated by zero-intelligence traders (no ability to maximize or even seek profits, or observe or remember market events), these markets are guaranteed to yield, on average, a high proportion of their surplus to the traders. A large part of the efficiency of these markets is the result of their structural properties, independent of the motivations or abilities of the traders who participate in them.
2. The apparently minor differences in the rules of continuously clearing and synchronized double auctions have important consequences for their surplus extraction properties. When compared to synchronized double auctions, the minimum expected efficiency of continuous double auctions is lower by about 6% while their mean efficiency over the feasible parameter range is lower by about 4%. It would be premature, however, to conclude on the basis of this result that the continuously clearing form is necessarily less desirable than the synchronized form of double auction in all situations. Continuous auctions may have superior price discovery properties and may therefore dominate synchronized markets in environments where the price discovery function is an important issue. The results and conjectures presented here only point to interesting directions for future work.
3. There are two possible causes of reduction in efficiency of auctions: (a) money left on the table by traders and (b) displacement of intra-marginal traders by extra-marginal traders. Money is left on the table when potentially profitable trades are unexploited at the end of the auction. Double auctions with inexperienced human traders often exhibit such behavior in early periods, but it tends to disappear with even a small amount of experience. The zero-intelligence machine traders of our markets repeatedly submit bids and asks from beginning of a period to the end; therefore, these markets do not suffer from this source of inefficiency. Displacement of intra-marginal traders is the only source of inefficiency in the markets examined here.
4. Efficiency of double auctions is influenced not only by the shape of demand and supply to the left of the equilibrium point but also by their shape to the right. However, the magnitude of this influence of extra-marginal units on expected efficiency is constrained by two countervailing forces. As the value of either α or β is increased from zero, chances that an efficiency-reducing transaction will occur increase; at the same time, the magnitude of efficiency reduction from such a transaction declines, yielding the efficiency-minimizing combination of parameters shown in Figure 2.
5. The ratio of extra- to intra-marginal traders in the market determines the magnitude of shortfall in expected efficiency of double auctions. As this ratio increases, expected efficiency declines as the extra-marginal traders increase their chances of displacing the intra-marginal traders. This effect is especially pronounced in continuously clearing double auctions.
6. Of the total surplus exploited in a synchronized double auction, a significant proportion (from 50–100%) is extracted within the first round itself, leaving only a small amount for the next few rounds, and little for the rest. This time

profile of surplus extraction measures the efficiency gains arising from the repeat bidding feature of double auction.

ZI traders can serve as a powerful tool for analysis and comparison of market institutions because they help isolate the consequences of market structure from the behavior of market participants. The behavior of human traders may change in response to changes in economic institutions, making it difficult to isolate the impact of changes in market structure alone on the basis of data from the field. The use of artificial traders is a convenient tool for holding behavior constant, while the structure of the economic institutions is varied to examine their consequences.

ACKNOWLEDGMENTS

We have benefitted from comments of Dan Friedman, Charles Plott, John Rust, Vernon L. Smith, Robert Wilson, and from the participants in presentations given at Carnegie-Mellon University, Economic Science Association, University of Kobe, and Yale University. Detailed comments of an anonymous referee, and from various participants at Santa Fe Institute and Bonn workshops were especially valuable. Financial support for this research was provided by Richard and Margaret Cyert Family Fund, Deloitte and Touche Foundation, National Science Foundation under grant SES-8915225, Santa Fe Institute, and Deutsche Forschungsgemeinschaft, Gottfried-Wilhelm-Leibniz-Forderpreis is gratefully acknowledged. We alone are responsible for any errors.

APPENDIX

DERIVATION OF EXPECTED EFFICIENCY OF SYNCHRONIZED DOUBLE AUCTION

Expected efficiency of a synchronized double auction populated by budget-constrained zero-intelligence traders (with one intra-marginal buyer and seller each, an infinity of extra-marginal buyers and sellers) is given by:

$$\begin{cases} 1-(1-\alpha-\beta)\left(\dfrac{\alpha^2}{1-\alpha(1-\alpha)}+\dfrac{\beta^2}{1-\beta(1-\beta)}\right) & \text{for } \alpha+\beta\leq 1\,, \\ 1 & \text{for } \alpha+\beta>1\,. \end{cases} \tag{4}$$

To derive this expression, let a_1 and b_1 be the ask and bid submitted by the intra-marginal seller and buyer respectively in the first round. Both a_1 and b_1 submitted by "zero intelligence" traders are distributed independently and uniformly over [0, 1]. In addition, the asks submitted by "zero intelligence" extra-marginal sellers are distributed uniformly over $[(1-\alpha), 1]$. As the number of extra-marginal sellers increases without bound, the lowest ask submitted by these sellers converges to $1-\alpha$. Bids submitted by "zero intelligence" extra-marginal buyers are distributed uniformly over $[0, \beta]$; as the number of extra-marginal buyers increases without bound, the highest bid submitted by these buyers converges to β.

In order to calculate the expected surplus extracted in this market when $(\alpha + \beta) < 1$, divide the square in (a_1, b_1) plane into nine cells as shown in Table 1.

CELL 1:. Both intra-marginal bid and ask are less than β. Probability of this event is β^2. The intra-marginal bid $b_1 < \beta$ is outbid by the maximum of the extra-marginal bids (at β). This highest extra-marginal bid of β is crossed with the intra-marginal ask $a_1 < \beta$ to effect a transaction and realize a surplus of β. (At this point we are concerned only with the total surplus, and not with how this surplus is split between buyers and sellers.) With $(\alpha + \beta) < 1$, extra-marginal buyers and sellers cannot trade with each other. Therefore the only remaining trade possible is between the intra-marginal buyer and an extra-marginal seller. Since the current ask submitted by these sellers is $(1-\alpha)$, it is only a matter of time when, in subsequent rounds of bidding, the intra-marginal buyer submits a bid higher than $(1-\alpha)$ and a transaction takes place, realizing a surplus of α. Thus total surplus of $(\alpha+\beta)$ is extracted in Cell 1, yielding an expected surplus of $\beta^2(\alpha+\beta)$. Since the maximum expected surplus that could have been extracted is β^2, it represents a loss of $g_1 = \beta^2(1-\alpha-\beta)$ in expected surplus.

TABLE 1 Intra-Marginal Asks and Bids

	Bid b_1 from intra-marginal buyer		
	Cell 7 Prob: $\alpha\beta$ $(1-\alpha) < a_1 < 1$ $0 < b_1 < \beta$	**Cell 8** Prob: $\alpha(1-\alpha-\beta)$ $(1-\alpha) < a_1 < 1$ $\beta < b_1 < (1-\alpha)$	**Cell 9** Prob: α^2 $(1-\alpha) < a_1 < 1$ $(1-\alpha) < b_1 < 1$
Ask a_1 from intra-marg. seller	**Cell 4** Prob: $\beta(1-\alpha-\beta)$ $\beta < a_1 < (1-\alpha)$ $0 < b_1 < \beta$	**Cell 5** Prob: $(1-\alpha-\beta)^2$ $\beta < a_1 < (1-\alpha)$ $\beta < b_1 < (1-\alpha)$	**Cell 6** Prob: $\alpha(1-\alpha-\beta)$ $\beta < a_1 < (1-\alpha)$ $(1-\alpha) < b_1 < 1$
	Cell 1 Prob: β^2 $0 < a_1 < \beta$ $0 < b_1 < \beta$	**Cell 2** Prob: $\beta(1-\alpha-\beta)$ $0 < a_1 < \beta$ $\beta < b_1 < (1-\alpha)$	**Cell 3** Prob: $\alpha\beta$ $0 < a_1 < \beta$ $(1-\alpha) < b_1 < 1$

CELL 9: This case is analogous to Cell 1; the expected surplus associated with this cell is $\alpha^2(\alpha+\beta)$, or a loss of $g_9 = \alpha^2(1-\alpha-\beta)$ in expected surplus.

CELL 2: In this cell, the intra-marginal buyer outbids the extra-marginal buyers ($b_1 > \beta$) and the former crosses with the intra-marginal ask ($\alpha_1 < \beta$), yielding a surplus of 1 with probability $\beta(1-\alpha-\beta)$. Expected surplus associated with this cell is therefore the maximum possible $\beta(1-\alpha-\beta)$ and $g_2 = 0$.

CELL 6: By argument analogous to Cell 2, the expected surplus associated with this cell is the maximum possible $\alpha(1-\alpha-\beta)$ and $g_6 = 0$.

CELL 3: Again, intra-marginal units transact with each other, yielding a surplus of 1 with probability $\alpha\beta$ or expected surplus of $\alpha\beta$ equal to its maximum possible value and $g_3 = 0$.

CELL 5: Since the intra-marginal buyer outbids the extra-marginal buyers in this cell, and intra-marginal sellers have outbid the extra-marginal sellers in the first round, the extra-marginals have no chance of entering trading. The bid/ask improvement rule of double auction makes it impossible for these traders to hold the current bid/ask in a subsequent round after the first round bids fall in this cell. Intra-marginal units will necessarily trade with each other (either in the first round if the intra-marginal bid exceeds the intra-marginal ask, or in a later round of bidding). In any case, there can be no loss of surplus once the intra-marginal bid and ask occupy this cell. Therefore, $g_5 = 0$.

When the first round bids and asks fall in Cells 4, 7, or 8, no transaction can take place until subsequent rounds of bidding.

CELL 7: This cell probability in the first round is $\alpha\beta$. Probabilities in the second and subsequent rounds in this cell are exactly the same as in the first. If g is the expected loss of surplus for this double auction, $\alpha\beta g$ is the expected loss of surplus associated with this cell.

CELL 4: The first round probability of this cell is $\beta(1-\alpha-\beta)$, and it forces a second round of bidding. If intra-marginal bids/asks submitted in the second round fall in Cells 2, 3, 5, or 6, no loss of surplus takes place. If the second-round submissions are in Cell 8 and 9, the maximum bid/minimum ask over the two rounds falls in Cell 5 and 6 respectively, again leading to transactions with no loss of surplus. If second-round submissions are in Cell 1 (probability β^2), surplus-reducing transactions cause an expected loss of $\beta^2(1-\alpha-\beta)$. If second-round submissions are in Cell 7 (probability $\alpha\beta$), the minimum ask/maximum bid fall in Cell 4. Thus, with probability $\alpha\beta + \beta(1-\alpha-\beta) = \beta(1-\beta)$, bidding is forced into a third round, and so on. Thus the probability of loss, from the first round bids falling in this cell, is given by:

$$\beta(1-\alpha-\beta)\beta^2(1+\beta(1-\beta)+\beta^2(1-\beta)^2+\beta^3(1-\beta)^3+\ldots) = \frac{(1-\alpha-\beta)\beta^3}{1-\beta(1-\beta)}.$$

The corresponding expected loss from first round bids falling in this cell is

$$g_4 = \frac{(1-\alpha-\beta)^2\beta^3}{1-\beta(1-\beta)}.$$

CELL 8: By argument analogous to Cell 4, the expected loss of surplus from Cell 8, g_8, is given by

$$g_8 = \frac{(1-\alpha-\beta)^2\alpha^3}{1-\alpha(1-\alpha)}.$$

Let g be the expected loss of surplus from all nine cells.

$$\begin{aligned} g &= \sum_{i=1}^{9} g_i \\ &= \alpha^2(1-\alpha-\beta) + \frac{\alpha^3(1-\alpha-\beta)^2}{1-\alpha(1-\alpha)} + \beta^2(1-\alpha-\beta) + \frac{\beta^3(1-\alpha-\beta)^2}{1-\beta(1-\beta)} + \alpha\beta \times g \cdot \\ &= (1-\alpha-\beta)\Big(\frac{\alpha^2}{1-\alpha(1-\alpha)} + \frac{\beta^2}{1-\beta(1-\beta)}\Big) \end{aligned}$$

Expected efficiency is

$$1-g = 1-(1-\alpha-\beta)\Big(\frac{\alpha^2}{1-\alpha(1-\alpha)} + \frac{\beta^2}{1-\beta(1-\beta)}\Big).$$

REFERENCES

1. Becker, Gary S. "Irrational Behavior and Economic Theory." *J. Pol. Econ.* **70** (1962): 1–13.
2. Easley, D., and John Ledyard. "Theories of Price Formation and Exchange in Double Oral Auctions." This volume.
3. Friedman Daniel. "On the Efficiency of Experimental Double Auction Markets." *Am. Econ. Rev.* **74**(1) (1984).
4. Friedman, Daniel. "The Double Auction Institution: A Survey." This volume.
5. Gode, Dhananjay, and S. Sunder. "Allocative Efficiency of Markets with Zero Intelligence (ZI) Traders: Market as a Partial Substitute for Individual Rationality." Carnegie-Mellon University Working Paper, 1991.
6. Gode, Dhananjay, and S. Sunder. "Human and Artificially Intelligent Traders in Computer Double Auctions." Presented at the meetings of Economic Science Association, Tucson, 1989.
7. Palmer, R., J. Rust, and J. Miller. "Participant's Manual for a Double Auction Tournament." Manuscript, Santa Fe Institute, 1989.
8. Plott, Charles R. "Industrial Organization Theory and Experimental Economics." *J. Econ. Lit.* **20** (1982): 1485–1527.
9. Rust, J., R. Palmer, and J. Miller. "A Double Auction Market for Computerized Traders." Working Paper 92-02-008, Santa Fe Institute, 1992.
10. Rust, J., J. Miller, and R. Palmer. "Behavior of Trading Automata in a Computerized Double Auction Market." This volume.
11. Satterthwaite, Mark A., and Steven R. Williams. "The Bayesian Theory of the k-Double Auction." This volume.
12. Smith, V. L. "An Experimental Study of Competitive Market Behavior." *J. Pol. Econ.* **70** (1962): 111–137.
13. Smith, V. L. "Experimental Economics: Induced Value Theory." *Am. Econ. Rev.* **66** (1962): 274–279.
14. Smith V. L. "Microeconomic Systems as an Experimental Science." *Am. Econ. Rev.* **72**(5) (1982).
15. Wilson, Robert. "On Equilibria of Bid-Ask Markets." Technical Report 452, Institute for Mathematical Studies in Social Sciences, Stanford University, 1984.