On the Efficiency of Experimental Double Auction Markets

By Daniel Friedman*

A double auction is a real world, do-it-yourself market procedure in which participants may make public offers both to buy ("bids") and to sell ("asks"). Experimentalists have for some time noted its "surprising" efficiency; for instance, after listing several aspects of this efficiency in his recent survey article for this *Review*, Vernon Smith remarks:

> Perhaps the most important general feature of the experimental results... is the support they provide for what might be termed the Hayek hypothesis: [Double Auction] *Markets economize on information in the sense that strict privacy together with the public messages of the market are sufficient to produce efficient* C[ompetitive] E[quilibrium] *outcomes.* [1982, pg. 947].

In his recent survey, Charles Plott comments, "The competitive model seems to work best when markets are organized as oral double auctions.... The overwhelming result is that these markets converge to the competitive equilibrium even with very few traders" (1982, p. 1493). Neither author attempts a theoretical explanation; indeed, at the end of his article, Smith leaves the "Hayekian" or competitive market performance as a "scientific mystery."

The goal of this paper is to solve the mystery, at least in part. Of particular interest is how agents, who may initially be quite ignorant of others' circumstances and strategies, may be guided by the double auction institutions to transact at market-clearing prices.

The issue is important for both theoretical and practical reasons. The Walrasian model that underlies most of our theorizing about competition suggests that a large number (strictly speaking, an uncountably infinite number) of agents, not the handful actually employed in experiments, are required to produce competitive market outcomes (compare Robert Aumann, 1964). Hence these experimental results raise doubts regarding the cogency of the Walrasian model. On the practical side, one notes that many of our most important financial markets, including most U.S. money markets, commodities markets, and interbank foreign exchange and Eurocurrency markets, already are conducted as double auctions. Moreover, the SEC's 1971 Institutional Investor Study recommended that the NYSE abandon its traditional "specialist" institution in favor of an electronic double auction, perhaps similar to that employed by NASDAQ in recent years. If there are sound theoretical reasons to believe that a double auction structure would enhance market efficiency as much in the stock market as it evidently has in laboratory experiments, economists would be able to make strong policy recommendations to the SEC for reform of the NYSE.

It is perhaps surprising in view of their prevalence in the world and in experiments that double auctions have so far received very little attention from theorists. David Easley and John Ledyard (in the only direct theoretical study of which I am aware) point out:

> A participant in a Double Oral Auction has a very complex decision problem. He must decide when to bid [or ask]; how much to bid [or ask] and, whether or not to accept the trades

*Visiting Assistant Professor of Economics, Schools of Business, University of California, Berkeley CA 94720 This work was partially supported by the Foundation for Research in Economics and Education, and by the Center for the Study of Futures Markets, Columbia University I thank Glenn Harrison and Jon Salmon, my collaborators in related experimental work, for helping me focus on the theoretical issues examined here. An early version of this paper was presented at a UCLA workshop, and participants gave useful suggestions I also owe a considerable debt to two referees for thoughtful criticisms that helped me recast and sharpen the exposition.

Copyright © 2001 All Rights Reserved

> offered by other subjects. Further, all of these decisions must be made with very imperfect information. The subject does not know the payoffs or expectations of other agents, he does not know the terms of trade that will be available to him in the future, and he does not know the effect of his actions on the actions of others. This is a very complex incomplete information game.... [1983, p. 8]

They circumvent many of these complexities by assuming that agents maintain "reservation prices" that adjust over time in plausible ways, and that agents follow bid, ask, and acceptance rules[1] based on their current reservation prices. For a slightly simplified version of a popular experimental market structure, they prove asymptotic convergence of transacted prices and final allocations after many repetitions of the market to (approximate) competitive equilibrium, at least for some parameter configurations. They also test some implications of their model against experimental data.

The results I derive here are generally complementary to those of Easley and Ledyard. I focus on the efficiency properties of a *single* repetition of the market, although I also discuss informally the asymptotic behavior. I take a more general view of the strategies and learning rules that may be employed by participants in a double auction, and rely on an apparently mild consistency condition, phrased in terms of a simple game in normal form, to generate the results. I find that the double auction will generate outcomes that are *ex post* Pareto efficient in some cases and nearly so in others if the consistency condition (called "No-Congestion Equilibrium") is satisfied.

A systematic analysis of market performance requires several steps. One must first characterize the market participants and market institutions. This step is relatively straightforward, and is taken at the beginning of the next section. Next one must derive tractable descriptions of individual actions that arise from the first step. Many approaches are possible here, and I will devote the remainder of the second section to a quick survey of them and to a motivation for my own approach. The third and final step, characterizing the market outcomes arising from the second step, is begun in a very simplified setting (the "basic case") in the third section. In Section IV, the analysis is extended to a setting commonly employed in experimental markets. The final section summarizes the insights arising from the analysis, suggests extensions, and discusses possible applications to experimental data and to contemporary financial markets.

[1]Subject to certain restrictions mentioned below in fn 5.

I. Market Institutions and Strategies

The market exists over a continuous finite time interval $[0, T]$. It is populated by a finite number of agents, often referred to as traders, indexed $i=1,\dots,n$. In the cases examined here, there are only two goods, denoted m and x. Good m is the numeraire, referred to as money or cash. The other good x is sometimes referred to as the asset. Each trader i is characterized by

1) *Endowment*, nonnegative holdings (m_i^0, x_i^0) of the two goods, not both zero, at time $t=0$;

2) *Preferences*, represented by a smooth monotone von Neumann-Morgenstern utility indicator U_i defined over *final* $(t=T)$ holdings of the two goods; and

3) *Information*, summarized at each time t by a set consisting of public information available to all traders at time t as well as private information available only to the given trader; these sets will be specified more fully below.

Our market is a *double auction*: at each time $t \in [0, T]$, each trader i sets a nonnegative *ask price* $p_i^a(t)$, representing the amount of cash he is prepared to accept in exchange for a single unit of the asset. At the same time, each trader sets a nonnegative *bid price* $p_i^b(t)$, representing the amount of cash he is prepared to part with to acquire a single unit of x. The "best" ask price $p_M^a(t)=\min_i p_i^a(t)$, and the "best" bid price $p_M^b(t)=\max_i p_i^b(t)$ are referred to as the *market* ask and bid, respectively. Trader i can buy a unit of x at

Copyright © 2001 All Rights Reserved

time t by accepting the market ask price; the resulting transaction consists of a unit increase in i's holding of x and a decrease of $p_M^a(t)$ in his cash holding, together with offsetting changes in the holdings of his counterparty, the trader j who held the market bid at t. The case of trader i selling a unit of x to the trader j who holds the market bid is analogous. Transactions that would yield negative holdings of money or the asset for either party are not allowed, and traders' bids and asks are restricted accordingly. Holdings $(x_i(t), m_i(t))$ adjust only through transactions occurring according to these rules.

It is assumed that new public information at time t consists only of the current market bid and ask prices, and an indication of whether a transaction occurred at either of these prices. Thus each trader at time t has available the (partial) history of these public prices, as well as that for his own prices and transactions. It is also assumed that each trader has private prior information, including at least knowledge of his own endowments and preferences.

It may help to have an example of such an institutional structure. Imagine a dozen widget traders scattered across town, each with a terminal linked to a central computer. Each terminal screen displays current market bid and ask prices (in dollars per widget), as well as the trader's own current bid and ask prices and current holdings. If he needs to refresh his memory, he can also get a display of past prices, his previous trades, and the time remaining in the trading day. Whenever he chooses to do so, a trader can push a button to accept the current market ask. The central computer would then verify that our trader has sufficient cash in his account, and then execute the transaction by updating its records of our trader's holdings of cash and widgets, as well as those of the counterparty, who held the market ask. Similarly, a trader can signal acceptance of a market bid price and the transaction will be consummated as long as he currently has at least one widget. A trader can also revise his current bid or ask prices by typing in the appropriate commands. Traders may not know the number and characteristics of other traders present in the market, nor the identity of their counterparties in transactions. At the end of the trading day, traders can pick up their final holdings of cash and widgets at a central warehouse.

I now turn to the derivation of individual actions. Implicit in the characterizations of traders is the following rationality postulate:

ASSUMPTION A: *Each trader's actions at each time t are chosen to maximize expected utility of final holdings, given information currently available to him, his current holdings, and the double auction rules.*

The actions referred to, of course, are the announcement of one's own current bid and ask prices, as well as the choice of which current market bid and ask prices will be accepted. I note in passing that a relatively simple rule for the latter choice is to pick a maximum acceptable market ask price, $v_i^b(t)$ (and a minimum acceptable market bid price, $v_i^a(t)$) and to buy (respectively, to sell) by accepting the market price whenever it is no greater than $v_i^b(t)$ (respectively, no less than $v_i^a(t)$). If such prices always exist and satisfy the coherency condition $p_i^a(t) \geq v_i^a(t) \geq v_i^b(t) \geq p_i^b(t)$, we say that i follows a *reservation price strategy*.[2]

The term strategy is used advisedly. The effect of one's own action depends crucially on the actions of others. For instance, announcing a bid of \$1.00 will have absolutely no effect if someone else currently maintains a bid of \$1.10, but in the event that everyone else is currently bidding less than \$1.00, it will become the market bid and may result in an immediate transaction. In the latter case, it will provide public information to other participants, which may affect their current expectations and subsequent actions, which in turn may well affect one's own future trading opportunities. Hence once we specify precisely how current information is incorporated into current expectations in Assumption A (a "learning rule"), we have a very

[2]My 1982 paper, Proposition A2, provides a sufficient condition for reservation price strategies to be consistent with the Assumption A. I will not use reservation price strategies in any essential way in this paper, but do refer to them on occasion.

Copyright © 2001 All Rights Reserved

complicated continuous-time extensive game of incomplete information.

Following the approach pioneered by John Harsanyi (1967; 1968a, b) and more recently articulated by David Kreps and Robert Wilson (1982) among others, one might try to analyze the double auction as follows. First, make decision times and actions discrete,[3] and postulate as common knowledge that traders all believe that others' parameters (preferences and endowments) are drawn from known distributions. Strategies may then be regarded as complete contingency plans, in which the action (or probability distribution over actions) selected by trader *i* at time *t* depends on the public and private information available to him at that time, including in particular the market price histories. One looks for equilibrium selections of strategies and beliefs, for which 1) beliefs (i.e., probability distributions regarding others' parameters, or more germanely, regarding the trading opportunities that others will make available) are updated when new information arrives via Bayes' Rule, using likelihood functions based on others' equilibrium strategies, and 2) strategies are optimal in the sense of Assumption A relative to equilibrium beliefs. Such noncooperative equilibria are variously called Perfect Bayes Equilibria or Sequential Equilibria, according to how one handles events of (equilibrium) probability zero.

This approach to our problem is not very promising for several reasons. The first is its inherent complexity. Drew Fudenberg and Jean Tirole (1983), for example, require a rather intricate analysis of this type even to analyze a market with only two traders, two discrete decision times, and extremely simple initial parameter distributions. One is not encouraged to attempt an extension to more traders, much less to something approaching continuous time. A different sort of simplification, arising from the work of Lloyd Shapley and Martin Shubik, can be found in Pradeep Dubey (1982). One collapses the auction to a game in *normal* form in which strategies consist of one-shot, price-quantity schedules of purchase and sales offers, and outcomes are based on a clearing mechanism. Dubey finds that certain ("active") Nash Equilibria of these games coincide with Competitive Equilibria for the underlying economy under quite general conditions. Leo Simon (1982) reaches a similar conclusion in a related context, emphasizing the role of "Bertrand competition" in forming equilibrium prices. Unfortunately, this work is based on a complete information concept of equilibrium that is inappropriate to our present concerns. Robert Wilson (1982) obtains some results (on *ex ante* efficiency) for a *sealed-bid* double auction using an incomplete information equilibrium concept, but it is not clear how his approach (based on the auction literature following William Vickrey, 1961) can cast any light on how publicly announced bids and offers can provide market participants with sufficient information to achieve outcomes that are efficient *ex post*.

Hence, even heroic simplifications[4] of the explicit incomplete information extended-form game approach to double auctions have so far shown little promise in resolving Smith's mystery. The basic problem is that, within this framework, any learning rule that is fully consistent with Assumption A seems to presuppose an incredible amount of prior information (as well as computational ability!). In the Perfect Bayes Equilibrium concept, for instance, each trader must a priori know, for any realization of parameters, what strategies others would follow. If one asks how such parameter-contingent strategies (themselves information-contingent actions) could possibly be learned if they are *not* known a priori, one is driven to postulate some Nash equilibrium in meta-learning rules, as in Larry Blume et al. (1982). Of

[3]For instance, one might suppose that a computer implementing the double auction has a known cycle time and sequence for polling all traders for acceptances or posted price changes. One can also restrict all prices to integer multiples of some minimal unit, say one cent or a mill

[4]Another type of simplification, suggested by Glenn Harrison, is to employ Assumption A in a Bayesian learning model in which traders ignore strategic considerations. We are currently investigating the explanatory power of this "Bayesian Games Against Nature" approach. See also Harrison, Smith, and Arlington Williams (1981)

Copyright © 2001 All Rights Reserved

course, this only pushes the analysis to a new level of complexity while begging the question of what meta-meta learning process could generate the given meta-equilibrium, and so on, *ad infinitum*.

I avoid the dilemma here by an indirect approach. Instead of postulating some explicit "rational" learning rule and a lot of prior information, I seek plausible restrictions on learning and strategies that will yield a tractable and insightful analysis of the double auction for minimal prior information. Of course, I face the danger that if the restrictions are not well chosen, I will obtain only uninteresting tantologies. The reader will have to judge for himself.

Perhaps the best way to begin is to note that there are only four actions that can immediately increase a trader's utility of *current* holdings:

1) accept the market bid ("sell at market") if it exceeds one's marginal valuation (based on current holdings) of the asset;

2) accept the market ask ("buy at market") if it is below one's current marginal valuation;

3) seize ("raise") the market bid if one does not already hold it and it is below one's current marginal valuation; and

4) seize ("shave") the market ask if one does not already hold it, and it is above one's current marginal valuation.

Actions 1) and 2) yield $\Delta U_i(t) > 0$ directly, while 3) and 4) lead to an expected improvement in one's current bundle as long as one believes that there is a positive probability (perhaps quite small) that someone will accept the current market ask or bid. We will say that an agent is *myopic*[5] if he always adopts one of these four actions whenever at least one is available.

However, myopia is not always the best strategy. After all, Assumption A says that traders attempt to increase current expected utility of *final* holdings, rather than utility of current holdings. A trader who believes he knows where prices are going might, for instance, accept a market ask price *above* his current marginal valuation, intending to resell later at a higher price, and thus more than recoup his "short-term loss." It is even possible that a trader may engage in such speculation in order to influence the beliefs of others, as is illustrated in a bit of folklore surrounding Wellington's defeat of Napoleon at Waterloo. The story goes that it was common knowledge in the London stock market on the day of the battle that Nathan Rothschild would be the first to learn the outcome. When Rothschild very publicly sold shares at fairly low prices, other participants inferred that Napoleon had won, and prices dropped precipitously as panic selling ensued. Rothschild's agents then snapped up the shares incognito, providing the wily financier with a large profit.[6]

A third type of behavior, perhaps the most likely if agents are initially quite ignorant of each others' characteristics, is simply to hold back—not accept favorable market prices and not seize the market price—in hopes of transacting at more favorable prices later in the period.

However, as the end of the trading period approaches, the opportunities for gain from holding back or from speculating decline. Traders are then increasingly likely to accept favorable market prices and/or seize the market price, since these are the only ways to actually realize gains in a double auction market. It is possible that some traders might wait too long before employing these myopic strategies and, in a last minute flurry of bids, offers and transactions, some potential gains might go unrealized. However, given any detailed specification of how bids and offers are processed, it seems plausible that traders who know these procedures a priori or have learned them would be able to avoid getting shut off from attempted transactions at the end of the trading round.[7] One does observe

[5]Easley and Ledyard in effect assume that agents always employ strategies that are "myopic" relative to current *reservation prices* rather than current marginal valuations. The NP-process of my (1979) article formalizes myopic trading in a more general context.

[6]This example also illustrates the point that rational traders need not employ reservation price strategies Rothschild's acceptance set for market bids must have included a range of prices that were too high to pass up as well as a disjoint range below recent transacted prices that would convey the intended misleading signal.

[7]Indeed, the danger of last-second market congestion clearly will induce rational traders to transact early, particularly those with larger potential gains from a single transaction (see Lorraine Glover, 1983), larger

Copyright © 2001 All Rights Reserved

in actual experiments that, although there is sometimes a flurry of transactions late in a trading round, experienced subjects seldom miss out on attempted transactions by waiting too long.

These considerations lead to my restriction on agents' learning rules and strategies, *no-congestion equilibrium*, a consistency condition that says in a rather weak sense that our agents are able to achieve a Nash equilibrium in their final actions. Roughly speaking, I ask: if the market were unexpectedly held open an extra instant, would anyone definitely wish to change his bid or ask prices, or accept the market bid or ask after all? If not, we have a no-congestion equilibrium.

One can formalize this idea in terms of a hypothetical *extra chance game*, a single instant, simultaneous-move (i.e., normal form) game appended to our double auction game. Strategies in the extra chance game consist of possibly resetting closing bid or ask prices (in particular shaving or raising the market) and possibly accepting the closing market bid or ask. A strategy of neither resetting prices nor accepting at the market is referred to as status quo. Payoffs for the extra chance game consist of changes in the U_i's generated by transactions according to double auction rules from final holdings in the double auction game. We say that beliefs are *nondogmatic* for the extra chance game if each trader believes that acceptance of the market bid by some trader has positive probability (perhaps quite small), and similarly for the market ask.

DEFINITION: *The strategies played in the double auction game constitute a no-congestion equilibrium, if, for some set of nondogmatic beliefs, the status quo strategy maximizes expected payoff for each player in the associated extra chance game.*

Note that this condition appears to be quite weak in that it only requires that there be *some* beliefs that support the closing actions in the double auction game as *ex ante* optimal; the agents' final actions need not constitute a Nash equilibrium *ex post*. However, it is clear that traders would play myopic strategies in an extra chance game, and this turns out to have significant consequences for market efficiency, as will be shown in the next two sections.

II. The Basic Case

Experimental markets typically involve indivisibilities of the asset and "corner solutions" (zero holdings of money or asset), analytical complications that may obscure the nature of our results. For that reason, I first examine in this section a very simple pure exchange economy, referred to as the basic case, that resembles the world studied in intermediate price theory. For the remainder of this section only, the following assumptions will be in force. With regard to agents' preferences, in addition to the maintained assumptions of smoothness and monotonicity, we now also assume convexity (i.e., $U_i^{-1}(c,\infty)\in\mathscr{R}_+^2$ is always strictly convex) and a boundary condition (i.e., $U_i^{-1}(c)$ doesn't intersect the axes). Cobb-Douglas preferences, for instance, satisfy all these assumptions. It is also assumed that initial endowments are all strictly positive, so corner solutions are no longer possible at any time. With regard to auction rules, we allow unrestricted "odd-lot" trading when a bid (or offer) has been accepted: each party to the transaction (say $i=1,2$) picks a $\lambda_i\in(0,1]$ representing the largest fraction of a unit he desires to transact. Then holdings of both parties change by the fraction $\lambda=\min\{\lambda_1,\lambda_2\}$ of the standard double auction transaction specified earlier in the previous section. Thus the asset is regarded as perfectly divisible.

It will be convenient to refer to i's marginal rate of substitution of the asset for money at current holdings, $g_i(t)=\partial_2U_i(m_i(t),x_i(t))/\partial_1U_i(m_i(t),x_i(t))$, where ∂_k denotes the kth partial derivative. Note that our maintained assumptions of smoothness and monotonicity ensure that $g_i(t)$ is well-defined and positive. Also note that U_i increases for a small purchase (sale) at price p iff $g_i>p$ (respectively, $g_i<p$), and that by convexity, g_i decreases (increases) for small purchases (sales) at prices near g_i.

desired volume, greater risk aversion, or certain types of asymmetric information.

Copyright © 2001 All Rights Reserved

We can now neatly characterize equilibrium states for our basic case. Eschewing the slightly ambiguous term "competitive equilibrium," we have a *price equilibrium* (Gerard Debreu, 1959, p. 93) in a (feasible) allocation (m^*, x^*) and a price p^*, if, for each $i = 1, \ldots, n$,

$$(1) \qquad g_i(m_i^*, x_i^*) = p^*;$$

that is, the allocation is Pareto optimal and supported by p^*. If we also have for each $i = 1, \ldots, n$,

$$(2) \qquad m_i^* + p^* x_i^* = m_i^0 + p^* x_i^0,$$

then we have a *Walrasian equilibrium*. Note that the budget condition (2) will generally be violated by trading at the "false prices," (i.e., $p(t) \neq p^*$) that can be expected early in a double auction.

PROPOSITION 1: *If there are $n \geq 3$ traders in a double auction market, and the strategy selection is consistent with Assumption A and results in a no-congestion equilibrium, then the outcome of trading constitutes a price equilibrium: the final holdings are Pareto optimal, and closing market bid and ask prices coincide and support the final holdings.*

PROOF:

A simple arbitrage argument eliminates the possibility that the market bid ever exceeds the market ask. Hence, if the closing market bid p_M^b and closing market ask p_M^a do not coincide, we may assume without loss of generality that these market prices are held by traders 1 and 2, respectively, and $p_1^a > p_2^b$. Letting g_3 denote trader 3's marginal rate of substitution at the final holdings, we have three cases. If $g_3 \geq p_1^a > p_2^b$, then trader 3 could increase his expected utility in an extra chance game either by accepting the market bid (buying from trader 1) until $g_3 \leq p_1^a$, or else (given nondogmatic priors) by seizing the market bit from trader 2. Either way trader 3 will not employ a status quo strategy, so this case is impossible in no-congestion equilibrium. Similarly, $p_1^a > p_2^b \geq g_3$ is inconsistent with no-congestion equilibrium: trader 3 will sell to trader 2 or shave trader 1's ask. But the remaining case $p_1^a > g_3 > p_2^b$ can't occur in no-congestion equilibrium either; trader 3 could gain by buying at a price higher than the market bid or selling at a price lower than the market ask, so he would either "shave" or "raise" the market prices. Therefore, $p_M^a = p_M^b$; call this common value p.

Suppose now that $g_i > p$ for some i. Clearly i does not hold the market bid and could profit by accepting it in an extra chance game until $g_i \leq p_M^b$, so the supposition is false in no-congestion equilibrium. Similarly $g_i < p$ is impossible in no-congestion equilibrium: i should sell at the market ask. Hence $0 < g_i = p$ for all i.

It is perhaps surprising that only three agents are required to yield efficient final allocations under the apparently weak assumption of no-congestion equilibrium. With only two agents, a bilateral monopoly impasse could occur in which buyer and seller both have overly optimistic beliefs and each looks to the other to make price concessions, even in an extra chance game. The presence of third agents breaks the impasse in a double auction market, because then the buyer (or seller) must compete to hold the market bid (or ask) or else be frozen out of the market. Thus, a form of "Bertrand competition" is at work.

A natural question to ask at this juncture is whether the outcome could actually constitute a Walrasian equilibrium. A sufficient condition for Walrasian equilibrium (in addition to the hypotheses of the proposition) is that all transaction prices coincide, but this condition can be ensured by an appropriate choice of reservation price strategies. Specifically, suppose that p^* is a Walrasian equilibrium price, the existence of which is well-known under much weaker conditions than those presently in force (for example, Debreu). If each trader i sets $p_i^a(t) = v_i^a(t) = \max\{p^*, g_i(t)\}$, and $p_i^b(t) = v_i^b(t) = \min\{p^*, g_i(t)\}$, then one can verify[8] that the outcome will be a Walrasian equilibrium and furthermore, that these reservation price

[8]See my (1982) paper, Proposition 3.2.

Copyright © 2001 All Rights Reserved

strategies constitute a (complete information) Nash equilibrium of the double auction game as well as a no-congestion equilibrium. This example incidentally establishes the consistency of the assumptions of the proposition.

Of course, such strategies are unlikely to be employed unless traders have a lot of consistent prior information. *Stationary repetition* as employed by experimenters since Smith (1962) is a device that provides just such information. The market is run several times, with each agent given exactly the same endowment and preferences in each repetition that he had in previous repetitions. The only difference across repetitions in agents' characteristics is that each has priors that are based on a larger public information set (viz, the price history of previous double auctions) in later repetitions.

Given a specific learning scheme (or suitable restrictions on such schemes), one might argue that under stationary repetition the outcome of a double auction, starting from an arbitrary price equilibrium in the initial repetition, will eventually evolve to a Walrasian equilibrium. For instance, if each trader follows some Bayesian procedure for incorporating public information into his (subjective) distributions of available buying and selling prices, then after a sufficient number of stationary repetitions these distributions will tend to converge, as long as no trader had dogmatic priors. The reason, of course, is that the public information (prices observed in previous repetitions) will eventually outweigh the (possibly diverse) first-repetition prior beliefs in the traders' later-repetition posterior distributions. A second, more delicate step in the argument is that the average subjective distributions (to which the individual distributions tend over repetitions of the market) eventually become concentrated at a point. The idea here is that as beliefs (i.e., probability distributions) converge, the range of transacted prices narrow, and vice versa. A final step, more routine, is to verify that this mutually reinforcing process is consistent only if a Walrasian price is the point of concentration. Unfortunately, this program is difficult to carry out in detail, and I will not attempt it here.

III. Experimental Markets

I now turn from the basic case to a specification of the double auction that is popular among experimentalists, in which the asset is indivisible, payoffs are separable in the asset and money, and traders are specialized as buyers or sellers. Hence in this section I will employ the following assumptions in addition to those of Section II. Preferences are of the form[9] $U_i(m,x)=m+h_i(x)$, where h_i is a nondecreasing function defined for x a nonnegative integer. An agent's endowment either satisfies $m_i^0 \geq h_i(1)$ and $x_i^0=0$, in which case i is called a *buyer*, or satisfies $m_i^0=0, x_i^0 \geq 1$, in which case i is called a *seller*.

For simplicity, I further specialize to the case of single unit transactors, that is, $x_i^0=1$ for sellers, and for each agent there is some $g_i>0$ such that $h_i(x)$ equals g_i if $x>0$, and equals 0 if $x=0$. (In this case, the marginal value of second unit of the asset is zero for everybody, so buyers have no incentive to acquire more than one unit of the asset.)

It is not hard to verify in this context that if a price p^* and allocation (m^*, x^*) satisfy

$$(3) \qquad x_i^* = 1 \text{ if } g_i > p^*,$$
$$x_i^* = 0 \text{ if } g_i < p^*, \quad i=1,\ldots,n,$$

then we have a price equilibrium. If we also have the usual budget condition

$$(4) \quad m_i^* + p^*x_i^* = m_i^0 + p^*x_i^0, \qquad i=1,\ldots,n,$$

then we have a Walrasian equilibrium. Henceforth, I shall for convenience index our agents so that $g_1 \geq g_2 \geq \cdots \geq g_n$. Then it is clear that Walrasian equilibrium prices p satisfy $g_k \geq p \geq g_{k+1}$, where $k=\Sigma x_i^0 \leq n$ is the total quantity supplied of the asset.

[9] In experiments, preferences are typically induced by paying subjects in "real money" for their final holdings of "cash" and asset according to some prearranged schedule that reduces to U_i or a linear transformation of U_i. (For instance, a fixed fee may be deducted from m_i.) Sometimes the rules also prohibit resale or repurchase of assets, a practice I will comment on at the end of this section.

Copyright © 2001 All Rights Reserved

TABLE 1: PARAMETERS FROM BOHM-BAWERK'S HORSE MARKET

Buyers		Sellers	
Agent Index	g_i	Agent Index	g_i
1	£30	4	£26
2	28	5	25
3	26	8	21.5
6	24	11	20
7	22	14	17
9	21	16	15
10	20	17	11
12	18	18	10
13	17		
15	15		

Note: $x_i^0(0)=0$ for buyers; $x_i^0(0)=1$ for sellers.

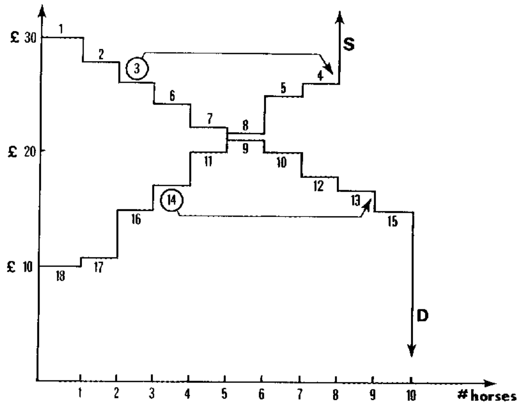

FIGURE 1. SUPPLY AND DEMAND FOR THE HORSE MARKET

A venerable example of such a market from Eugen von Bohm-Bawerk (1891) is reproduced in Table 1, with agents re-indexed according to my scheme. The assets ("horses") are to be traded against cash, denominated in pounds, in a market whose rules Bohm-Bawerk does not fully specify.[10]

Figure 1 displays the staircase-shaped supply and demand curves for this example; the labels attached to the steps indicate the corresponding agent indices. The Walrasian prices consist of the interval [21, 21.5] for which the two curves coincide. One can think of a transaction as reshuffling of the steps—for example, if 3 buys from 14, then as indicated by the arrows in Figure 1, step 3 on the D curve moves out to a position of the same height on the S curve, and likewise, step 14 of S becomes part of the D curve, while the gaps in their former positions are closed by sliding the outer portions of the staircases leftward. A transaction can occur to mutual advantage as long as at least one step of S remains below its corresponding D step. Hence we see that condition (1) above for Pareto efficiency amounts to saying that at the final allocation, the S and D curves do not intersect in the interior of the first quadrant.

What sort of outcomes might we expect from a double auction in this sort of market? Certainly transacted prices $p(t)$ should be bounded above by g_1 and below by g_n, and we might expect that buyers with low numbers and sellers with high numbers would be more likely to transact early. To the extent that Bertrand competition is effective, the results from the previous section would lead us to expect that the closing prices and allocations would approximate a "price equilibrium." However, after most transactions have occurred, the crucial "third trader" of Proposition 1 may not be so readily available here. For instance, suppose in the numerical example that early transactions gave rise to the configuration illustrated in Figure 2a. Agent 7, the highest valuation buyer remaining, can be expected to hold the market bid, but he can hold it as low as £21 and still outbid 9, his closest rival. Likewise 11 might well hold the market ask as high as $g_8 =$ £21.5. If agents 7 and 11 each thought the probability of a last minute acceptance by the other sufficiently great, an impasse could arise. Impasse seems more likely in the configuration[11] depicted in Figure 2b: agent 8 could hold the market bid as low as £20 and

[10] In the first few paragraphs of his argument (p. 204), Bohm-Bawerk seems to be describing an oral double auction. However, in view of his subsequent discussion, a more plausible interpretation is that he envisions an ascending multiple unit auction (in which buyers drop out when the bid price rises above their valuations), together with a descending auction for sellers. He neglects the possibility of "non-truth-telling" strategies. In any case, Bohm-Bawerk predicts a Walrasian outcome.

[11] This configuration could arise if 8 sold (to 1, say) and 9 bought (from 17, say) in early trading.

Copyright © 2001 All Rights Reserved

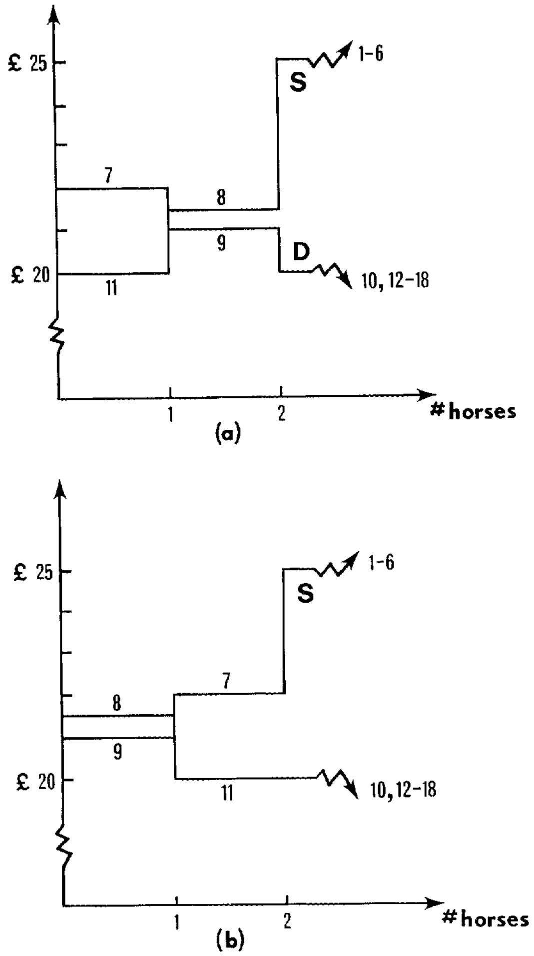

FIGURE 2. INEFFICIENT FINAL CONFIGURATIONS

9 could hold the market ask as high as £22, in which case neither could gain by accepting the market price.

It turns out that such local impasses, involving the four agents that Bohm-Bawerk calls the "marginal pairs," are the worst departure from price equilibrium one can expect in this setting. In making this claim more precise in the next proposition, I will call an allocation (m, x) *almost Pareto optimal* if at most a single transaction is required to make it Pareto optimal. Finally, note that due to ill-conceived speculation in early trading (i.e., buying high and selling low), a buyer could be effectively frozen out of the market by a shortage of cash at the end of trading—say, $x_i(T)=0$ and $m_i(T)<p<g_i$. This possibility presents no substantial difficulties to this analysis, but it does make my results much more awkward to state.[12] Therefore, it will be assumed for notational convenience that buyers are never cash constrained (i.e., $m_i(t) \geq g_i$ if $x_i(t)=0$).

PROPOSITION 2: *If the number of sellers $k \geq 2$ and the number of buyers $n-k \geq 2$, and if traders employ a no-congestion equilibrium strategy selection consistent with Assumption A, then closing market prices satisfy $g_{k-1} \geq p^a_M(T) \geq p^b_M(T) \geq g_{k+2}$, and the final allocation $(m(T), x(T))$ is almost Pareto optimal.*

PROOF:

Let $J_s = \{i | x_i(T)=1\}$ and $J_b = \{i | x_i(T)=0\}$ denote, respectively, potential sellers and buyers at final holdings. Given nondogmatic priors, each seller in an extra chance game would shave the market ask price if it is sufficiently high, so in no-congestion equilibrium, the closing market ask must be held by the seller i_s with lowest g_i, at a price at or below the second lowest $g_i \in J_s$. For similar reasons, the closing market bid price must be at or above the second highest $g_i \in J_b$, and be held by i_b, the buyer with highest $g_i \in J_b$. Hence

$$\min{}_2\{g_i : i \in J_s\} \geq p^a_{i_s}(T)$$
$$= p^a_M(T) \geq p^b_M(T)$$
$$= p^b_{i_b}(T) \geq \max{}_2\{g_i : i \in J_b\},$$

that is, the second smallest g_i in J_s is at least as large as the second largest g_i in J_b. Now $J_s = k$ by definition, so $\min_2\{g_i : i \in J_s\}$ is g_{k-1} if $i_b \leq k$, and is g_k otherwise. Similarly, $\max_2\{g_i : i \in J_b\}$ is either g_{k+1} or g_{k+2}, so the claimed bounds on closing market prices follow. If $i_b \geq k+1$ then the closing allocation is actually Pareto optimal by criterion

[12]In essence, one has to re-index buyers to reflect "effective" valuations $\min\{m_i(t), g_i\}$, rather than "notional" valuations g_i. See my paper (1982, Section 4) for a treatment of this complication in a related context

Copyright © 2001 All Rights Reserved

(3) above. Otherwise we must have $i_s \geq k+1$. In this case, if trader i_b were to purchase from trader i_s, then their roles would be reversed (i.e., then $i_b \geq k+1$), so such configurations are almost Pareto optimal; indeed, the single transaction required to achieve Pareto efficiency is that which generates the smallest mutual gains from trade.

The minimal sort of bilateral impasse possible in almost Pareto optimal allocation is not inevitable; by a slight variant of the argument employed after Proposition 1 we can establish the possibility of a Walrasian equilibrium outcome. The only changes necessary are to set a trader's ask price to a prohibitively high level (say $B = \Sigma_i m_i^0$) when he has no units of the asset, and to require that a trader's bid price never exceed his current cash balance.

The learning requirements for achieving something close to Walrasian equilibrium in stationary repetition of an experimental double auction actually seem less severe than in the basic case. For example, it seems plausible in a given repetition that a trader would refuse to buy at prices above the closing ask, or sell at a price below the closing bid, of the previous repetition. In the present context, such policies are quite consistent with no-congestion equilibrium strategies, and enable traders to improve upon their least favorable transactions without reducing trading volume. If all traders follow such policies, there will be a definite tendency for transacted prices to lie in narrower intervals in later repetitions. One presumes that the possible bargaining problem between i_b and i_s would eventually be resolved to mutual satisfaction, in which case final allocations in subsequent repetitions would be Walrasian, and transacted prices would lie in the Walrasian interval $[g_{k+1}, g_k]$. Again, I will not attempt to make this argument rigorous here.

To summarize, for the simple case of a popular experimental design considered in this section, I have established theoretically that closing prices and allocations will be almost efficient even in the first repetition of an experimental double auction market, and have suggested that they will ultimately converge to Walrasian equilibrium values under stationary repetition. These conclusions did not require that traders initially possess any information on others' circumstances or that they employ sophisticated learning rules. Neither did they require that agents be numerous: two traders on each side of the market suffice. Thus I have provided a partial theoretical explanation of the empirical regularities noted in Smith's (1982) Propositions 4, 5, and 6.

I close this section with a few caveats. If agents may transact more than one unit of the asset, the theoretical results are not quite so strong. For some parameter values, a "marginal" agent may have the incentive to act as a local monopolist, and outcomes may not then be quite Walrasian.[13] However, for parameter values commonly chosen in actual experiments, there are no such incentives, and my conclusions remain valid. If repurchases and resales of the asset are prohibited (see fn. 9), then the bounds on closing prices in Proposition 2 may be somewhat looser. In terms of a *S-D* diagram such as Figure 1, transactions may be represented in this case by *removing* the transactors' steps from the diagram, rather than reinserting them to the right. As a result, the "final configurations" as in Figure 2 may have second steps of the *S* and *D* curves (representing the closing price bounds) that are further apart. The indexing is more complicated in this case, but the reasoning is unchanged.

A final caveat is that experimental subjects are human, and may blunder or find ways of colluding. For reasons that should be clear from this analysis, these phenomena should quickly become less important in a double auction market as the number of subjects increases beyond the minimal two on each side of the market. Likewise, the assumption of no-congestion equilibrium seems increasingly plausible, as traders gain experience in a double auction market.

IV. Discussion

The results of the previous two sections illustrate the power of three institutional fea-

[13] See my paper (1982, Appendix) for an example of this phenomenon in a related context.

Copyright © 2001 All Rights Reserved

tures of experimental markets in shaping efficient outcomes. The primary feature is the double auction structure: in order to realize any gains from exchange, a trader must either seize the market price, or accept the market price of another. This necessity both limits each agent's influence on prices (via Bertrand competition) and also conveys high quality information to other agents. By contrast, under the "tatonnement" institution implicit in most discussions of price formation (in which each agent provides an excess demand schedule to an outside "auctioneer," who then computes and announces a market-clearing price), agents may be able to influence prices significantly by conveying misleading information (i.e., false excess demands) unless each agent is quite "small" relative to the size of the market.

The second feature is the preset time T at which trade will end. As the deadline T approaches, agents must become more myopic if they wish to realize any remaining gains from trade, thus enhancing the informational and competitive aspects of the double auction. Indeed, Proposition 1 suggests that these two features alone may suffice to bring about an efficient final allocation (or almost efficient, for Proposition 2) even if agents have very little prior information about each other.

The last feature is stationary replication, a device that makes agents' prior information both more consistent and more precise. I argued informally that under stationary replication, agents need not be very intelligent to learn enough to achieve a Walrasian outcome, for which all trades occur at a single price that clears the market. By contrast, such learning may be impossible in tatonnement markets, as argued in a different context by Roman Frydman (1982), for instance.

There seem to be no obstacles to extending arguments of the sort employed in Sections III and IV to more general cases. For instance, my 1982 paper examines a constant marginal rate of substitution specification employed[14] in experimental asset markets. The case of uncertainty (state-contingent utility functions) is analyzed in an Appendix to my paper with Harrison and Salmon (1983). The extension of the basic case of Section III to a many-asset market seems quite straightforward, at least if preferences are convex. One could easily justify an appropriate convexity assumption under (for instance) a mean-variance approach to asset valuation in which agents do not change their beliefs about the joint distribution of long-run asset returns during the trading process. Direct analogues of my present results appear valid in this case, although (due to lack of a boundary condition) one must qualify the three-trader assumption of Proposition 1 by three *active* traders, that is, for each asset, there are three traders who have positive holdings of cash and the given asset. Of course, when odd-lot trading is not permitted, one can't generally obtain equalized marginal rates of substitution, and the results are more awkward to state in this case. Nevertheless, it seems reasonable to conjecture that, once one has defined Pareto optimal and Walrasian equilibrium with appropriate indivisibility constraints, my results would still go through.

There are other ways to follow up the work presented here. First, it suggests new ways to analyze experimental data, in particular, the unaccepted bids and asks that have usually been ignored (Easley and Ledyard being a notable exception). One could postulate simple specific learning rules—for example, no learning (myopia), or the Bayesian Game Against Nature rules of Harrison, Smith, and Williams. The sequences of observable market bids, and asks and trades arising from Assumption A and reasonable priors could then be generated and compared to the experimental data.

Finally, we may glean new insights regarding price formation in contemporary financial markets. Although the trading procedures of such markets have increasingly come to resemble those postulated in my double auction model, one hardly expects stationary repetition in ongoing markets; to the contrary, new outside information (both public and private) typically arrives *during* the trading process. The end of a trading day, T, is

[14]See Robert Forsythe, Thomas Palfrey, and Plott (1982), for instance.

Copyright © 2001 All Rights Reserved

also usually much less important in this context, because participants are generally concerned with payoffs over a longer time span. Hence our results (extended perhaps to the multigood case) do *not* demonstrate that financial markets are efficient; indeed, to the extent that our conditions are violated, they may suggest reasons for the surprising degree of volatility in such markets.

REFERENCES

Aumann, Robert, "Markets with a Continuum of Traders," *Econometrica*, January 1964, *32*, 39–50.

Blume, L. E., Bray, M. M. and Easley, D., "Introduction to the Stability of Rational Expectations Equilibria," *Journal of Economic Theory*, April 1982, *26*, 313–17.

Bohm-Bawerk, Eugen von, *Positive Theory of Capital*, 1891; translated by W. Smart, Steckert: New York, 1923 edition.

Debreu, Gerard, *Theory of Value*, Cowles Monograph No. 17, New Haven: Yale University Press, 1959.

Dubey, Pradeep, "Price-Quantity Strategic Market Games," *Econometrica*, January 1982, *50*, 111–26.

Easley, David, and Ledyard, John, "A Theory of Price Formation and Exchange in Oral Auctions," Department of Economics Working Paper No. 249, Cornell University, rev. January 1983.

Forsythe, Robert, Palfrey, Thomas and Plott, Charles, "Asset Valuation in an Experimental Market," *Econometrica*, May 1982, *50*, 537–68.

Friedman, Daniel, "Money-Mediated Disequilibrium Processes in a Pure Exchange Economy," *Journal of Mathematical Economics*, March 1979, *6*, 149–67.

______, "Price Formation in Double Auction Markets," Working Paper No. 278, University of California-Los Angeles, November 1982.

______, **Harrison, Glenn and Salmon, Jon,** "The Informational Efficiency of Spot and Futures Markets," CSFM Working Paper, Columbia Business School, June 1983.

Frydman, Roman, "Towards an Understanding of Market Processes," *American Economic Review*, September 1982, *72*, 652–68.

Fudenberg, Drew and Tirole, Jean, "Sequential Bargaining with Incomplete Information," *Review of Economic Studies*, April 1983, *50*, 221–47.

Glover, Lorraine, "Microdynamic Price Determination," doctoral thesis proposal, University of California-Los Angeles, 1983.

Harris, M. and Townsend, R. M., "Resource Allocation Under Asymmetric Information," *Econometrica*, January 1981, *49*, 33–64.

Harrison, Glenn, Smith, Vernon and Williams, Arlington, "Learning Behavior in Experimental Auction Markets," University of California-Los Angeles, November 1981.

Harsanyi, John, "Games with Incomplete Information Played by Bayesian Players," Part I, *Management Science*, November 1967, *14*, 159–82.

______, (1968a) "...," Part II, *Management Science*, January 1968, *14*, 320–34.

______, (1968b) "...," Part III, *Management Science*, March 1968, *14*, 486–502.

Kreps, David and Wilson, Robert, "Sequential Equilibria," *Econometrica*, July 1982, *50*, 863–94.

Plott, Charles R., "Industrial Organization Theory and Experimental Economics," *Journal of Economic Literature*, December 1982, *20*, 1485–527.

Simon, Leo, "Bertrand and Walras Equilibrium," Working Paper, University of California-Berkeley, 1982.

Smith, Vernon L., "An Experimental Study of Competitive Market Behavior," *Journal of Political Economy*, April 1962, *70*, 111–37.

______, "Microeconomic Systems as an Experimental Science," *American Economic Review*, December 1982, *72*, 923–55.

Vickrey, William, "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, March 1961, *16*, 8–37.

Wilson, Robert, "Double Auctions," unpublished manuscript, Stanford University, 1982.

Copyright © 2001 All Rights Reserved

Copyright of American Economic Review is the property of American Economic Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.