Competitivity in Auction Markets: An Experimental and Theoretical Investigation

Daniel Friedman; Joseph Ostroy

The Economic Journal, Vol. 105, No. 428 (Jan., 1995), 22-53.

Stable URL:
http://links.jstor.org/sici?sici=0013-0133%28199501%29105%3A428%3C22%3ACIAMAE%3E2.0.CO%3B2-M

The Economic Journal is currently published by Royal Economic Society.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at http://www.jstor.org/about/terms.html. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at http://www.jstor.org/journals/res.html.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

http://www.jstor.org/
Wed Dec 22 19:01:10 2004

The Economic Journal, **105** (*January*), 22–53. © Royal Economic Society 1995. Published by Blackwell Publishers, 108 Cowley Road, Oxford OX4 1JF, UK and 238 Main Street, Cambridge, MA 02142, USA.

COMPETITIVITY IN AUCTION MARKETS: AN EXPERIMENTAL AND THEORETICAL INVESTIGATION*

Daniel Friedman and Joseph Ostroy

We report successive rounds of theory and laboratory experiment investigating price-taking behaviour and market efficiency. We focus on the impact of structural parameters as well as trading institution. The structural parameters involve non-competitive supply and demand, and a new odd-lot trading procedure for divisible goods. The trading institutions include the continuous double auction and the one-shot clearinghouse as well as a new quantities-only clearinghouse (CHQ) institution. We present and justify an as-if complete information theory which explains the competitive outcomes and which correctly predicts highly non-competitive outcomes in CHQ markets.

A finding from laboratory experiments challenges the conventional wisdom that competitive equilibrium arises only from large numbers of traders. Small numbers, say three or four buyers and three or four sellers, suffice to achieve essentially competitive equilibrium outcomes in repeated double auction market experiments (Smith, 1962, 1982*a*, *b*; Plott, 1982). Theoretical investigations by Wilson (1987), Friedman (1984, 1991*a*), Easley and Ledyard (1993), Satterthwaite and Williams (1989), and others have shed some light on this phenomenon, but the laboratory results remain puzzling.

This paper began as a sharp disagreement between the two authors as to the proper explanation for the puzzle. We investigate three approaches to reconciling theory and experiment.

(1) The traditionalist approach, initially favoured by one author, accepts the thrust of conventional wisdom, but notes that the number of traders is not the sole determinant of competitivity. There are special parameter configurations for which perfect competition – as defined by the inability of individuals to influence price favourably – is compatible with small numbers. It turns out that close approximations or even exact versions of these special configurations have been selected in most previous double auction experiments. Moreover, in virtually all previous experiments, subjects trade one or at most a few units of an indivisible good. Indivisibility greatly reduces profitable opportunities to influence price by withholding trade; indeed, a single-unit-trader can influence price in this manner only by withdrawing from the market and thus forgoing all profit. Holt *et al.* (1986) chose parameters intended to provide some traders with market power; even using indivisible units, they observed substantial departures from price-taking equilibria in some cases. More recently, van

* We have benefitted from the comments of Glenn Harrison, Charles Holt, Leo Simon and Vernon Smith, of session participants at the 1988 PCS/ESA meetings in San Francisco, and of seminar participants at UCLA and the University of New Mexico. We are indebted to the National Science Foundation and the UCSC Social Science Division for financial support, to Christopher McRae and Wyn Hillier for programming support and to Chanyong Park, Bret Carthew and Alexander von Borries for research assistance.

Boening and Wilcox (1994) observe sizeable departures from Pareto optimality in markets with discontinuous supply curves (arising from avoidable production costs) and no competitive equilibrium. Thus some existing evidence seems to support the traditionalist position that competitivity depends primarily on the parametric structure of the economy.

(2) The institutionalist approach initially favoured by the other author rejects the conventional wisdom and argues that the way the market is organised exerts a major influence on market outcomes. In particular, the dynamic nature of the double auction market institution (DA, also known as the continuous or oral double auction or as the bid–ask market) may somehow enhance competitivity. In the DA, every trader has the right to make public offers to buy (bids) or to sell (asks) and the right to accept others' offers at any moment during the trading period. A static market institution such as the Clearinghouse (CH, also known as the call market or the sealed double auction or sealed bid-offer market) might be less effective in forcing traders to act competitively. In the CH, bids and asks are not executed immediately but rather are collected during the trading period and at the end are aggregated respectively into demand and supply schedules which are cleared at a unified price.[1]

(3) A third approach, which we will call the as-if complete information Nash equilibrium approach (or complete information for short) occurred to us only after looking at the results of some experiments. The complete information approach has roots in the argument of Bertrand (1883) and suggests that price competition among even a few traders can produce competitive outcomes. The key theorems, due to Dubey (1982), Benassy (1986) and Simon (1984, 1987), establish a tight correspondence between the Nash equilibria of market games and competitive equilibrium outcomes. We shall argue that although traders' information in the experiments is far from complete, it is possible for them to learn to use the relevant 'complete information' strategies.

The first group of experiments we report provides a sharp test of the traditionalist approach by relaxing the structural constraints it identifies as 'forcing' competitive behaviour. The experiments employ (*a*) parameter configurations for demand and supply schedules chosen to maximise the gains to non-competitive behaviour and, equally important, they employ (*b*) a new DA market procedure called odd-lot trading that permits almost perfectly divisible units to be traded. In tandem, these changes in experimental design ensure that a single trader can have a large favourable influence on price by withholding a small amount of his competitive equilibrium trade. Nevertheless, and contrary to the traditionalist position, these experiments still produce highly competitive outcomes.

The next group of experiments tests the institutionalist view by using the same 'non-competitive' supply and demand schedules and odd-lot procedures

[1] Some economists might predict high efficiency in the CH but not in some other trading institutions. Although this may be a reasonable institutionalist position in some broad sense, we will use the 'institutionalist' label rather narrowly in this paper to refer to the prediction that the DA is a more efficient trading institution than the CH.

© Royal Economic Society 1995

in markets organised by the static CH institution. Nevertheless, and contrary to the institutionalist position, these experiments produce outcomes almost as competitive and efficient as those in the DA. Because they were explicitly formulated to produce the opposite, these DA and CH experiments yield some of the most convincing evidence to date of the robustness of competitive outcomes in auction market experiments.

In the attempt to understand why this happened, we re-examined the data and discovered that even when the true (experimenter-induced) supply and demand are very 'non-competitive', the supply and demand that traders choose to reveal to the CH typically are very 'competitive' – very elastic at market clearing prices and therefore not conducive to price or quantity manipulation. This discovery led us to formulate the complete information approach which (*ex post*) correctly predicts competitive outcomes in the CH (and by extension in the DA). These outcomes are the result of Nash equilibrium trader strategies that fully reveal quantity but that misrepresent induced reservation prices as being completely elastic at a market clearing price.

To test the predictive power of the complete information approach *ex ante*, we designed a new institution, called CHQ (quantity-only clearing house), that prevents traders from misrepresenting price and allows them only to choose a desired quantity. In theory, the complete information equilibrium of CHQ is that competitive equilibrium outcomes will not be observed. Consistent with this prediction, outcomes were very inefficient in the CHQ experiments.

In the next section we sketch a theoretical framework for studying the Clearinghouse in simple supply–demand economies. The implications of incomplete and complete information game theory models are summarised. In Section II, we present the experimental design and the hypotheses to be tested. The experimental results are collected in Section III. In the last section, we review and interpret our findings, point to some remaining questions and disagreements, and suggest directions for further work.

I. THEORETICAL CONSIDERATIONS

In this section we describe a class of general equilibrium models related to laboratory experiments. The models are fairly simple in that the only goods are a divisible 'money' and a traded good whose marginal rate of substitution is piecewise constant for each individual. We shall examine three versions of non-cooperative equilibrium in these models – non-manipulable, Bayesian Nash, and as-if complete information Nash equilibrium – as candidates both for prediction and explanation of experimental outcomes.

I.1. *The Model*

Let $I = B \cup S$ be the set of individuals, classified either as buyers or sellers. Each $i \in I$ is defined by two k-vectors, $(r_{i1}, \ldots, r_{ik})$ and $(q_{i1}, \ldots, q_{ik})$, indicating

© Royal Economic Society 1995

the reservation values and the quantities to which they are applicable. For each $i \in B$, reservation values decline with quantity in stair-step fashion, i.e. $r_{i1} \geqslant \cdots \geqslant r_{ik} \geqslant 0$ and $0 \leqslant q_{i1} \leqslant \cdots \leqslant q_{ik}$: r_{ij} is the reservation value per unit in the interval $(q_{ij-1}, q_{ij}]$, $j = 1, \ldots, k$ and $q_{i0} = 0$. For each $i \in S$, the situation is reversed: $0 \leqslant r_{i1} \leqslant \cdots \leqslant r_{ik}$, and $q_{ik} \leqslant \cdots \leqslant q_{i1} \leqslant 0$ and $q_{i0} = 0$, i.e. cost of supply is increasing in stair-step with quantity.

For i with parameters (r_{ij}) and (q_{ij}), suppose $y_i \in (q_{ij-1}, q_{ij}]$ for some $1 \leqslant j \leqslant k$. Evaluate the utility of y_i as

$$v_i(y_i) = \sum_{l=1}^{j-1} r_{il}(q_{il} - q_{il-1}) + r_{ij}(y_i - q_{ij-1});$$

and $v_i(0) = 0$. Fig. 1 illustrates Buyers' preferences; the utility v_i is the area under the staircase to the left of y_i. Throughout we shall assume that y_i is of the same sign as q_{ik} and $|y_i| \leqslant |q_{ik}|$. It is convenient to describe supply as negative demand, but in the Figs. below both quantities demanded and supplied will be represented as positive.

Let V_i^{2k} be the set of all such utility functions with at most $2k$ parameters. An economy is therefore defined by a

$$v \in V^{2k} = \times_i V_i^{2k},$$

where each i is chosen to be in B or S.

We call attention to two special cases. The most frequently studied class, both theoretically and experimentally, has $k = 1$ and $q_{i1} \equiv 1$. Call this class V^1 because each individual is defined by a single parameter, his/her reservation value, and because the bid/ask value is always for one unit of the commodity. This is distinguished from the unrestricted case of $k = 1$, called V^2, where both the reservation value and the quantity traded are free to vary. Since quantities are fixed in V^1, when we consider strategic manipulation of equilibrium below it will only be possible to falsify reservation values; however, in V^2 quantity withholding via odd-lot trading will also become a strategic possibility. We classify environments by V^h where $h = 1, 2, 4, 6, \ldots$. Note that if $h < h'$, then $V^h \subset V^{h'}$.

A competitive equilibrium, or more precisely, a Walrasian equilibrium (WE) for $v \in V^h$ consists of a price p and quantities (y_i) such that

(1) (utility maximisation) for each i, $v_i(y_i) - py_i \geqslant v_i(x_i) - px_i$, where $0 \leqslant x_i \leqslant q_{ik}$ if $i \in B$ and $0 \geqslant x_i \geqslant q_{ik}$ if $i \in S$.

(2) (market clearance) $\Sigma y_i = 0$.

The total payoff to the individual is the utility received from the traded commodity, $v_i(y_i)$, minus the amount of money paid out (if $y_i > 0$) or plus the quantity of money received (if $y_i < 0$).

Denote by $P(v)$ the set of prices p and by $Y(v)$ the set of allocations (y_i) which form part of some WE for v. Because preferences of individuals have the quasi-linear form, it is well known that the set of pairs of prices and allocations that constitute a WE for v is $P(v) \times Y(v)$.

© Royal Economic Society 1995

The Clearinghouse (CH) market institution operates as follows: trades are based on a Walrasian mechanism, by which we mean that after each individual announces the limit quantities and the limit prices at which he/she is willing

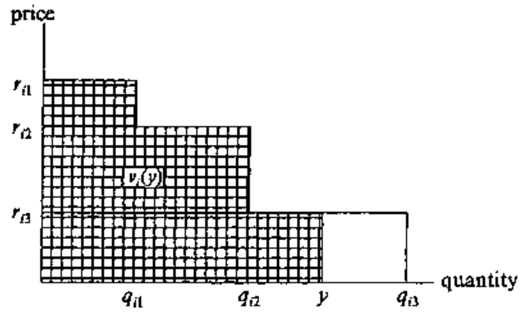

Fig. 1. Individual parameters and surplus.

to trade, the market mechanism collects that information to produce a Walrasian allocation. Formally, for a given h, if bids and offers are $w \in V^h$, the market aggregates these orders to produce an allocation $(p, (y_i)) \in P(w) \times Y(w)$.

An individual with characteristics $v_i = ((r_{ij}), (q_{ij})) \in V_i^h$ is free to submit any bid/offer $w_i = ((\hat{r}_{ij}), (\hat{q}_{ij}))$ provided it is feasible. This means first of all that $w_i \in V_i^h$ and also that $|\hat{q}_{ij}| \leqslant |q_{ij}|$. The latter restriction precludes an individual from making bids/offers exceeding his/her actual capacity, but it permits an individual to trade less (to withhold). There are no restrictions on reservation values other than those already defined. Call $V_i^h(v_i) \subset V_i^h$ the set of feasible announcements at v_i.

I.2. *Incomplete Information Equilibria and Inefficiency*

Suppose $v \in V^h$ and $(p(v), y(v))$ is a WE for v. Say that this equilibrium is non-manipulable if for all i and all $w_i \in V_i^h(v_i)$,

$$v_i[y_i(v)] - p(v)\, y_i(v) \geqslant v_i[y(v^i, w_i)] - p(v^i, w_i)\, y(v^i, w_i),$$

where v^i is v with v_i deleted and (v^i, w_i) is v with w_i in place of v_i. Otherwise, the WE is manipulable. In words, a WE is manipulable if some trader can submit a feasible bid or offer that (given truthful submissions by others) advantageously influences the allocation or price.

Non-manipulable equilibria are best known when applied to a class of economies. For example, a Walrasian mechanism is non-manipulable on $U \subset V^h$ if it is non-manipulable for each $v \in U$. In this case, we would say that truth-telling is a dominant strategy for a Walrasian mechanism on U.

Clearly, if WE were perfectly competitive in the everyday sense that no individual could influence prices, it would be non-manipulable; but we are dealing with a small numbers environment where individuals will typically be able to influence prices. We give a characterisation of perfect competition with small numbers.

© Royal Economic Society 1995

Denote by $V_*^h \subset V^h$ those v satisfying the following: there exists $p \in P(v)$ such that for each i, $p \in P(v^i)$, where $P(v^i)$ is the set of WE prices for the economy defined by v^i. See Fig. 2 for an illustration. Such a v has the property that

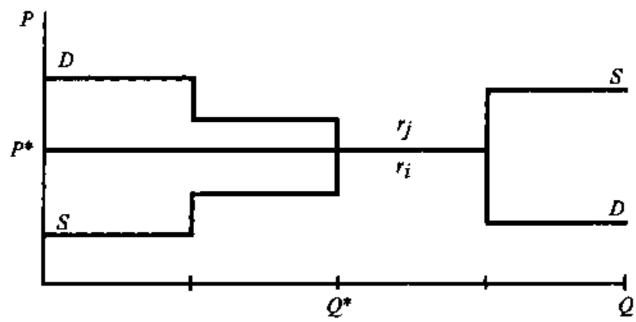

Fig. 2. A V_*^1 economy.

equilibrium prices do not depend on the participation of any one individual. Since a WE has the property that each individual's payoff is completely defined by prices, in V_*^h the utility each member of the group of all-but-one individuals can enjoy on his/her own is the same as each one receives when the remaining individual is added. Thus, each individual contributes no surplus to the others. The following two Propositions are implied by results in Makowski and Ostroy (1987).

Proposition 1. *For any h, if $v \in V_*^h$, then v is non-manipulable. Conversely, if $U \subset V^h$ is a convex set (in the space of parameters) and $U \subset V_i^h$ includes the constant functions $r_{i1} = r_{i2} = \cdots = r_{ik} = \alpha \geqslant 0$, then the Walrasian mechanism is non-manipulable only if $U \subset V_*^h$.*

When $h = 1$, the connection between perfect competition and non-manipulability is easy to characterise.

Proposition 2. *If $U \subset V^h$ is convex, then a necessary condition for $v \in V^h$ to be non-manipulable is that $P(v)$ be a singleton. If $h = 1$, then $P(v)$ is a singleton if and only if $v \in V_*^1$.*

Many laboratory economies restrict each buyer and each seller to single indivisible units, and so are V^1 economies. Most such economies discussed in the classic surveys of Plott (1982) and Smith (1982*b*) have substantial quantities with reservation values or costs at (or near to) a unique market clearing price. These are (or are near to) V_*^1 economies, so from a theoretical point of view it should not be surprising that such markets produce Walrasian outcomes.

Of course, V_*^1 is a thin subset of V^1, so in general the Walrasian mechanism is manipulable. As noted in the Introduction, the indivisibility of the non-money commodity in V^1 implies that the trader who attempts to influence price by 'shading' his/her reservation value risks being priced out of the market. In

© Royal Economic Society 1995

V^2 (or any other $h > 1$), the trader can withhold some but not all quantity in the attempt to create a favourable price change without being subjected to that risk. In some cases (see below), even small amounts withheld can have a markedly favourable influence on price. Therefore, one might expect that by designing an experiment in which v is chosen 'far' from the set of non-manipulable equilibria and the traded commodity is almost completely divisible, the outcomes of laboratory experiments would not be Walrasian.

These remarks lead us to the parametric configuration $v^{\square} \in V^2$, illustrated in Fig. 3, as the polar opposite of an element of V^1_*. In a box economy the buyers

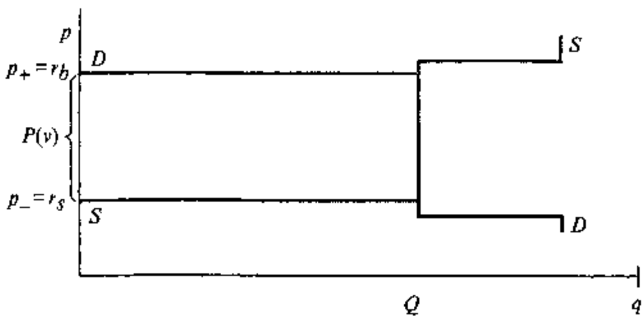

Fig. 3. A 'box' economy.

have the same reservation value r_b for a total of Q units of the traded good, and sellers have the same cost $r_s < r_b$ for the same total number of units. Notice that $P(v^{\square})$ is not unique; indeed the interval of WE prices is as wide as it can be given the trade volume Q and the total gains from trade for the economy as a whole. (Total gains are measured by the area between the demand and supply schedules to the left of the equilibrium quantity.) The wider the interval $P(v)$, the greater the potential profit from manipulating its WE.

Consider the strategic problem in $v^{\square}$. Assume for concreteness that the market mechanism chooses the mid-point of $P(v)$ as the market clearing price.[2] Then by withholding a very small quantity, the buyer (seller) can create a large and favourable change in price from the mid-point to the bottom (top) of $P(v)$. All participants can withhold in the attempt to reap the favourable price advantages of being on the short side of the market, so one might except that, in a box economy, trade and efficiency will be far less than in Walrasian equilibrium. However, this was not what we observed in experiments.

Therefore other equilibrium concepts may be worth considering. To begin, suppose there is a probability measure μ with full support on V^h (assumed to be a compact and convex set in the space of parameters) describing the frequency with which actual economies are chosen. Denote by μ^i the marginal on $\times_{j \neq i} V_j^h$ and assume that it is independent of v_i. The measure μ^i represents i's (correct) beliefs about the frequency with which characteristics of the rest

[2] The thrust of the following remarks do not depend on this selection.

© Royal Economic Society 1995

of the population are drawn. A (pure) strategy by player i is a mapping $\sigma_i: V^h \to V^h$ indicating the limit quantities and reservation values i will announce as a function of his actual characteristics, i.e. $\sigma_i(v_i) = w_i$. The strategy is feasible if $\sigma_i(v_i) \in V_i^h(v_i)$.

Given the strategy choices by others, $\sigma^i = (\sigma_j)_{j \neq i}$, the expected gain to v_i from announcing w_i is

$$\mathrm{E}_{\sigma^i}(w_i; v_i) = \int (v_i\{y_i[\sigma^i(v^i), w_i]\} - p[\sigma^i(v^i), w_i]\, y_i[\sigma^i(v^i), w_i])\, d\mu^i(v^i).$$

Given V^h and $\{\mu^i\}$, the strategy profile $\sigma(v) = (\sigma_i(v_i))$ is a Bayesian Nash Equilibrium (BNE) if for all i, v_i and $w_i \in V_i^h(v_i)$,

$$\mathrm{E}_{\sigma^i}[\sigma_i(v_i); v_i] \geqslant \mathrm{E}_{\sigma^i}(w_i; v_i).$$

As an important special case, suppose each σ_i is the identity mapping, i.e. the truth-telling strategy. Clearly, it is only when truth-telling is a BNE that the market mechanism will be efficient; otherwise, the economy will incur losses from non-competitive behaviour in the form of missed opportunities for trade due to underrevelation of reservation values[3] and/or reduced volume of trade due to strategic withholding.

It is well known that truth-telling is not a BNE for small numbers of traders. (See Myerson and Satterthwaite (1983) for the two trader version of V^1.) Qualitatively, the explanation is similar to that for manipulability. Results in Makowski and Ostroy (1989) imply that:

PROPOSITION 3. *Under the hypotheses of Proposition 1 applied to the support of μ, the* BNE *of a Walrasian mechanism will be truth-telling if and only if the support of μ is V_*^h.*

Quantitatively, however, the question arises as to how far from full efficiency a BNE is. If the discrepancy is small, this might explain the experimental results. For example, Satterthwaite and Williams (1993) report that expected efficiency losses decline rapidly ($O(n^{-2})$) as the number n of buyers and sellers increases. We point to three reasons why this finding does not provide an adequate explanation of previous and current experiments. The first is theoretical and the others have to do with the gaps between the laboratory observations and the BNE explanation.

First, note that the relative efficiency of BNE for small numbers is obtained with respect to the environment V^1. It follows from Proposition 2 that the prospects for favourable manipulation depend on the μ measure of the set $\{v: |P(v)| > \alpha\}$, the economies with Walrasian price intervals wider than α. The smaller the measure of this set for any given α, the smaller the expected gains for any strategy based on underrevelation of reservation values. With Lesbegue measure on V^h, $h > 1$, this set will be of null measure. Nevertheless, there are gains to withholding quantity to increase the chances of being on the short side

[3] Here and below, we refer to buyers' or sellers' understatements of the range of acceptable prices as 'underrevelation' This shorthand terminology is straightforward for buyers, but for sellers it means that announced willingness to accept is above true willingness (cost).

© Royal Economic Society 1995

of the market. As a theoretical proposition, the Satterthwaite–Williams results for V^1 do not carry over to V^2 and other economies with divisible traded goods.

Second, there are participants' beliefs. The gains to manipulation depend upon the properties of the perceived μ^i. Current experiments (and virtually all previous experiments) do not place traders randomly in an economy according to some diffuse measure. On the contrary, the current experiments were designed so that the actual economy is far from the competitive parameter configuration, with the implication that if others are truthfully revealing their characteristics, there are large and consistent gains to manipulation. If a measure concentrated on economies far from V^1_* had been used in the theory of BNE, the gains from manipulation and the resulting efficiency losses would have increased.

Last and perhaps most important, there is a gap between the logic underlying the BNE for a Walrasian mechanism and the behaviour observed in the experiments. The logic points to approximate efficiency because traders have rather small incentives to underreveal prices, whereas in the experimental outcomes efficiency is accompanied by significant price underrevelation as we shall document below.

To summarise the results of this section, careful examination of a class of general equilibrium models reveals that competitivity, as defined by non-manipulability of a CH (i.e. Walrasian) mechanism either in the deterministic or Bayesian setting, is compatible with small numbers. By designing experiments that lie as far as possible outside this competitive class, we set up a stronger test than heretofore of the robustness of Walrasian outcomes with small numbers.

I.3. *As-if Complete Information and Efficiency*

The gap between theory and experiment encouraged us to re-examine the data and this in turn led to a revision of the theory. We discovered that when actual parameters looked like D and S in Fig. 4 (the box-like configuration), the

Fig. 4. Actual and announced supply and demand. Announced economy (D′, S′) is V^2_*, but actual economy (D, S) is V^2.

© Royal Economic Society 1995

reported parameters looked more like D' and S', i.e. the strategic response of individuals with non-competitive parameters ($v \in V^2 \backslash V^2_*$) was to announce a pattern of bids and offers consistent with perfect competition ($w \in V^2_*$). We shall see that such non-truthtelling, but nevertheless perfectly competitive, reports are mutual best responses under fairly general conditions.

The equilibrium concept relevant to these observations is: $w \in V^h$ is a Nash equilibrium (NE) for $v \in V^h$ if for each i, $w_i \in V^h_i(v_i)$, and for all $\hat{w}_i \in V^h_i(v_i)$,

$$v_i[y_i(w)] - p(w)\, y_i(w) \geqslant v_i[y_i(w^i, \hat{w}_i)] - p(w^i, \hat{w}_i)\, y_i(w^i, \hat{w}_i).$$

To isolate the trivial case, notice that if each individual were to announce a w_i in which $q_{ik} = 0$ that would be a NE. This is to be contrasted with what we call the set of positive trade NE.

It readily follows that if $v \in V^h_*$, then announcing v ('truthtelling') is a NE. But there are many other ways to obtain NE. By contrast, the non-manipulability and the BNE concepts admit an element of $P(v) \times Y(v)$ as an equilibrium for v only when $v \in V^h_*$. After proving a theoretical result characterising the broad range of NE in V^2 economies, we became aware of the very general results of Dubey (1982), Simon (1984) and Benassy (1986). The key condition for efficient Nash equilibrium turns out to be that the economy permits active competition from both sides of the market: those $v \in V^h$ for which $(y_i) \in Y(v)$ implies that $\#\{i : y_i \neq 0,\ i \in B\} \geqslant 2$ and $\#\{i : y_i \neq 0,\ i \in S\} \geqslant 2$.

The statement of the following Theorem has been specialised to suit our purposes; the cited versions apply to more general models with many goods and preferences that are not necessarily quasi-linear.

DUBEY–SIMON–BENASSY THEOREM. *For any economy $v \in V^h$ permitting active competition from both sides of the market, the set of positive trade Nash equilibrium outcomes for the Clearinghouse market is precisely the set of Walrasian equilibria.*

The conclusion is very strong: except for economies with a natural monopolist or monopsonist, the CH trading institution can implement every competitive (WE) outcome as a Nash equilibrium (NE). Moreover, every NE of the CH market is either trivial (no trade) or competitive. Thus there is a very close correspondence between efficiency and non-cooperative equilibrium in CH markets.

The logic behind the DSB Theorem has roots in Bertrand's criticism of Cournot (1883) and is suggested by the key condition that there be at least two active buyers and two active sellers. Even when $v \notin V^h_*$, strategic 'mis'-representation leads traders to impose a perfectly competitive environment on themselves, i.e. they are led to choose some $w \in V^h_*$. Further, the outcome of this mis-representation is benign: an equilibrium w produces $(p, (y_i))$ with the property that

$$(p, (y_i)) \in P(w) \times Y(w) \cap P(v) \times Y(v).$$

The sellers must report low enough ask prices to 'meet the competition' at the clearing price $p (\hat{r}_{ij} \leqslant p)$ or else forgo all trading profit. In the absence of quantity constraints (rationing), reporting a price below p will either have no effect (if enough other sellers price exactly at p) or else will reduce profit.

© Royal Economic Society 1995

Similarly, buyers must report bid prices $\hat{r}_{ij} \geq p$ in order to transact, and can do no better than to report $\hat{r}_{ij} = p$. In this case, withholding quantity becomes unprofitable for the same reason as in V^h_*: it cannot affect price, so can only reduce profit. If strategies of this sort allow all traders to transact desired quantities then p must be a WE price. On the other hand, if the clearing price is not Walrasian then the strategies are not a NE because a quantity-constrained trader's best response is to price a bit more aggressively. The result is that WE outcomes in the Clearinghouse correspond to NE strategies which are truthtelling in the quantity dimension but strongly underrevealing in the price dimension.[4]

Some interpretive remarks are in order before proceeding. To implement the prescribed strategies, traders must be able to condition their bids and offers on market price. It is important to observe that such behaviour is not possible in a (one-shot) game of incomplete information – the setting for non-manipulable and Bayesian Nash equilibrium – because the (unknown) price is obtained only after bids and offers are announced. Formally, such conditioning is possible only in a game of complete information.

Our conclusion thus appears to lead to a paradox. Hayek (1945) is frequently cited to emphasise that market price produces a kind of minimally sufficient statistic of the divided knowledge, e.g. on individual preferences and endowments, that is not and never will be common knowledge. Hence, for an explanation of market efficiency to rely on a theory in which participants are assumed to know each others' characteristics seems to contradict this well-accepted feature of markets. We shall argue that our reliance on complete information is more terminological than substantive.

For the general class of games, Nash equilibria depend on the characteristics (payoffs) and the strategies of the players in complicated and delicate ways. Thus, for an arbitrarily given specification of strategic interactions, if the characteristics of the players were changed it would be difficult to trace out their consequences for equilibrium. So, in general, it may be presumed that to find a NE the players need to know the complete description of the game. This is not the case for the market games described above.

An important property of the CH 'game', not shared by games in general, is that it constructs market price as an essential summary statistic of the announced bids and offers. Relatedly, there are many strategic choices $w \in V^h$ which, when aggregated by the market institution, produce the same outcome: in particular, if the true economy is v, there are many w such that

[4] The DSB Theorem is actually three different models and sets of conclusions by three different authors. None of the models completely includes our Clearinghouse market, but the overlap among their results and our own is so substantial that here we omit an independent demonstration (available on request from the second author). The argument that every WE is a NE is more or less the same for each model: every individual truthfully reports his quantity but (under-)reports his reservation value as identical to a WE price; this is a NE. The argument underlying the converse is easier to establish in our Clearinghouse than in the other three models. Indeed, using an analog of the Clearinghouse in his model, Dubey (1982, p. 124) finds an example of a NE that is not a WE, which leads him to adopt another trading rule. When such an example is adapted to our environment, the problem disappears. The reason is that for Dubey buyers must rely on sales of other non-money commodities to pay for purchases, and this can be a binding constraint, whereas buyers in our model can use the money commodity to bid for purchases and are therefore unconstrained.

© Royal Economic Society 1995

$P(w) \times Y(w) \cap P(v) \times Y(v) \neq \emptyset$. If $(p,(y_i)) \in P(v) \times Y(v)$, then shading reservation values to lie near p while leaving quantities unchanged leads to the same outcome as submitting one's true characteristics. Further, there is a kind of 'natural' advantage to this kind of misrepresentation: by not revealing how much one would gain from trade while nevertheless signalling an interest in trading, the individual discourages efforts by others to manipulate price. Equilibria achieved in this way are compatible with incomplete information. All that can be inferred is a weak inequality that the reservation values and quantities of a trader must have been such that he/she did not lose from the exchange; but the amount actually gained remains private information.

Most laboratory markets are repeated. Repetition, rather than complete information, permits conditioning on previous observation of market price. Because price conditioning is possible in the experiment but information is not complete, we call this an as-if complete information explanation. Further discussion of this matter is taken up in Section IV.1.

Our as-if complete information theory (summarised in the DSB Theorem) came after observing efficient outcomes in the laboratory data, so we felt obliged to find a trading institution for which the theory predicts inefficient outcomes. We gain confidence in a theory if it can predict *ex ante* and if it can predict inefficiency as well as efficiency.

Price underrevelation is the key to efficiency in the DSB Theorem, so we designed a trading institution which enforces truth-telling in the price dimension. The quantity-only Clearinghouse (CHQ) works as follows. Price determination and transactions are exactly as in the basic CH, and messages (parameter reports) are also the same except that only quantity reports are allowed, with the price understood to be the true value. That is, for $v_i = ((r_{ij}), (q_{ij})) \in V_i^h$, trader i's message to the CHQ institution is of the form $w_i = ((r_{ij}), (\hat{q}_{ij}))$, where only $\hat{q}_{ij}$ is chosen freely, subject to the feasibility constraint on quantities. Denote the restricted set of strategies in CHQ by $VQ_i^h(v_i)$. Recall that V_i^1 denotes the single parameter strategy set indicating the announced bid/ask value of trader i's single indivisible unit of the traded good. Similarly, $VQ_i^2(v_i)$ is a single parameter strategy space, where if $v_i = (r_i, q_i)$, q_i need not be unity.

The appropriate definition of equilibrium is: w is a NE of $v \in V^h$ for CHQ if for each i, $w_i \in VQ_i^h(v_i)$ and for all $\hat{w}_i \in VQ_i^h(v_i)$,

$$v_i[y_i(w)] - p(w)\, y_i(w) \geqslant v_i[y_i(w^i, \hat{w}_i)] - p(w^i, \hat{w}_i)\, y_i(w^i, \hat{w}_i).$$

Whereas all positive trade NE are WE for the CH market, provided there are two active traders on each side of the market, under CHQ the results can be quite different.

PROPOSITION 4. *If $v \in V_*^h$, then truth-telling is a NE for CHQ. In contrast, suppose that $v \in V^2$ and $P(v)$ is not a singleton; therefore, by Proposition 2, $v \notin V_*^2$. Then there is no positive trade NE of v for CHQ.*

The proof is straightforward and omitted. The idea is simply that at any positive trade outcome of the CHQ institution, some buyer or seller will be able

© Royal Economic Society 1995

to influence price favourably by reducing the quantity offered slightly.[5] This proposition gives us the *a priori* prediction that efficiency and trading volume will be much lower in the CHQ than in the CH trading institution for box economies and other economies far from V^2_*.

II. EXPERIMENTAL DESIGN

Each of our 10 experiments employs experienced and profit-motivated undergraduates (usually eight at a time) to serve as traders in a repeated market (usually of 16–20 trading periods). In this section we describe the trading institutions in more detail, and explain how the odd-lot procedure allows divisibility of the traded good. Then we summarise the design details for each experiment and spell out the testable predictions. The interested reader will find useful background information in Copeland and Friedman (1987) and can obtain complete instructions on request from the first author.

II.1. *Market Institutions*

The standard market institution in experimental markets is the Double Auction (DA): at each moment during a trading period, each trader can offer to buy a single unit (a bid) and/or offer to sell a single unit (an ask), and at least the best such offers (lowest ask, highest bid) are publicly displayed. Traders can accept the offers, goods or cash inventories permitting, to produce an immediate transaction.[6] Some of our experiments employ a computerised version of the DA that allows trade of fractional units, as described in the next subsection. Traders are told how long each trading period will last (typically three minutes in a DA experiment) and that between periods their (r, q) parameters and their rôles as buyers or sellers may occasionally change.

Most of our other experiments employ a computerised Clearinghouse (CH) market institution. In our implementation, traders enter one or more bids or asks at their terminals as in the DA, but these offers are interpreted as limit orders and are not executed immediately. Rather, at the end of the trading period the bids and asks are aggregated respectively into market demand and supply curves, and the market is cleared in the usual fashion. That is, the price (or the midpoint of the interval of prices) is found at which the supply announced in the asks equals the demand announced in the bids, and all higher bids and lower asks are filled at this Walrasian price.

We also employ two alternative versions of the CH institution. In the first alternative, CHS, traders are allowed to enter only a single bid or ask. That is, buyers can bid for any chosen quantity q' at any single chosen positive reservation price r', (subject to the liquidity constraint that $r'q'$ does not exceed endowed cash) and sellers can choose an r' and an offered quantity

[5] The formalism in this section assumes traders employ only pure strategies. There may be mixed strategy NE of the CHQ that on average produce positive trading volume.

[6] These general features of the DA institution allow many variants, such as oral *vs.* computerised; specialisation of agents as buyers or as sellers *vs.* all agents allowed to both buy and sell; improvement rules for new bids and asks *vs.* free-form, etc. We spell out the variants used in the present experiments in the next subsection.

© Royal Economic Society 1995

q' not exceeding endowment q. In terms of the notation of Section 1, $w_i = (r', q') \in V_i^2(v_i)$. In the second alternative, CHQ, each trader can specify only a single quantity q, with the price automatically set at the true reservation value, r_{i1} or r_{i2}. That is, each trader's announced supply or demand is restricted to be a chosen initial portion of his true (induced) demand or supply schedule, i.e. $w_i \in VQ_i^h(v_i)$.[7]

II.2. *Divisibility and Induced Valuation*

For reasons explained in Section 1, all the experiments reported here employ a device which allows trade of a divisible good. The default bid or ask involves a quantity called a 'round lot'. In our Double Auction experiments, traders may request a fraction from 0·01 to 1·00 of a round lot when entering or accepting a bid or ask, and the actual transaction is the smaller of the fraction requested by the buyer and the fraction requested by the seller. We refer to such a fractional unit (and to the device itself) as an 'odd-lot'. Thus odd-lot transactions can be any integral percentage of the round lot, for practical purposes a close approximation of perfect divisibility.[8] Similarly in the Clearinghouse and its variants, traders can request integral percentages (hundredths) of a round lot in each bid and ask. We sometimes refer to hundredths of a round lot as 'shares'.

A second, less important design innovation allows traders to participate on both sides of the market (e.g. buy and then resell) as in asset market experiments, but without the usual asset market restriction of a single constant marginal valuation. We specify two cutoff points, e.g. $q_{i1} = 50$ and $q_{i2} = 200$ shares, with corresponding marginal valuations (reservation prices or costs), e.g. $r_{i1} = \$1.10$ and $r_{i2} = \$0.75$. The rule is that traders can redeem the first q_{i1} shares held at the end of the trading round for r_{i1} per share and remaining shares up to q_{i2} at r_{i2} per share; shares held in excess of q_{i2} are worthless. In the numerical example, 250 shares (or 2·5 round lots) held at the end of the trading period would be worth (\$1.10) (50) + (\$0.75) (150) + (\$0.00) (50) = \$167.50 in accounting dollars. To summarise, each trader's endowed characteristics can be described by a vector $(r_{i1}, r_{i2}; q_{i1}, q_{i2}) = v_i \in V_i^4$.

Half of our traders (the 'buyers') are endowed with no shares but enough cash to prevent the liquidity constraint from binding under normal circumstances. Their excess demand thus coincides with their demand, as indicated in Panel A of Fig. 5. The other traders (the 'sellers') are endowed with no cash but a quantity of shares $\bar{q}$ (endow) typically between q_{i1} (limit 1) and q_{i2} (limit 2). Their excess demand can be represented by a supply curve as in Panel B of Fig. 5. At any price exceeding r_{i1} (pay 2), the trader depicted there

[7] The CH institution has direct counterparts in contemporary financial markets (see Cohen *et al.* 1986). The CHS institution is feasible as long as one can prevent traders from acting through surrogates. The CHQ institution, however, is not feasible unless true value and cost parameters are known to the Clearinghouse. Recent laboratory experiments at Arizona and at UCSC have investigated other variants of the CH institution which provide information on the order flow; see McCabe *et al.* (1993) and Friedman (1993) for example.

[8] The closest precedent in the experimental literature is Gray and Plott (1987). They allow trade of up to four (minimum size) units in a single transaction in their otherwise conventional oral (noncomputerised) double auction market.

© Royal Economic Society 1995

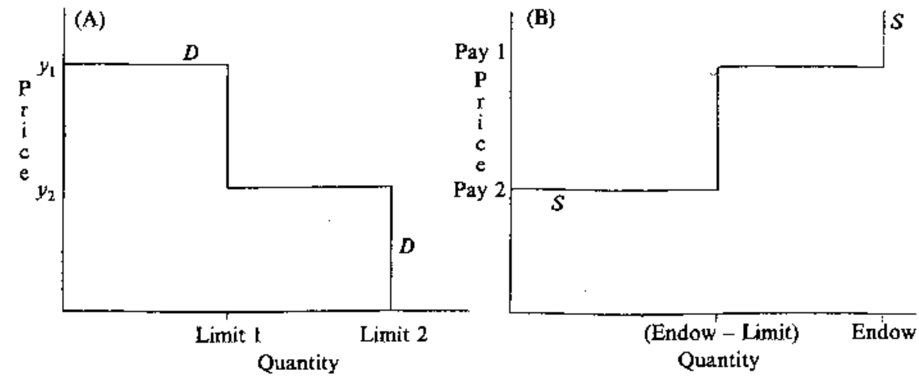

Fig. 5. Induced demand (A) and supply (B).

finds it profitable to sell his first $(\bar{q}-q_{i1})$ shares, leaving an inventory of q_{i1} shares; the remaining shares can be sold profitably only at a price exceeding r_{i2} (pay 1).

Each trader's profit in each trading period is the familiar buyer or seller surplus, measured in the usual way by differences between reservation values (r_{i1} or r_{i2}) and transaction prices, accumulated over transacted quantities. At the end of the experiment we paid each trader in US currency a fixed fraction (usually 0·01) of his or her profit in accounting dollars for all trading periods. The average trader went home with about \$20 at the end of a 90 to 120 minute experiment; a few traders earned \$35–\$40.

II.3. *Design of Experiments*

Our experiments are intended to compare market institutions given structural parameters that encourage price manipulation. For reasons discussed in Section I.1, we chose supply/demand schedules of a box economy (as in Fig. 3) and perturbations of it. Box (or box-like) economies have previously been investigated experimentally by Smith and Williams (1989) and by Holt *et al.* (1986), but only in the context of indivisible unit Double Auction markets.

Fig. 6 shows the r_{ij} and q_{ij} parameters used in our experiments for each buyer (D) and seller (S). Schedule IIa is the only 'competitive' induced economy we used; when its schedules are aggregated they lead to a $v \in V_*^h$. (Its unique WE price of \$0.65 still presents a difficult environment for competitive behaviour because most gains from trade in WE go to buyers.) Schedule IIb defines a 'box' economy with supply equal to demand over the price interval [\$0.65, \$1.75]. The other schedules define small perturbations of a 'box' economy – each trader can produce a box configuration by unilaterally announcing less than his or her true quantity. For example, in a schedule Ia economy with 4 sellers, the announced economy would have a WE price interval of [\$0.50, \$1.60] if a single seller revealed 140 shares, an underrevelation of $4 \times 20 = 80$ of her 220 shares. Schedule Ib presents exactly similar opportunities for buyers, who receive no gains given full revelation. Buyers and sellers have symmetric opportunities in schedule III economies; if each buyer (or each seller)

© Royal Economic Society 1995

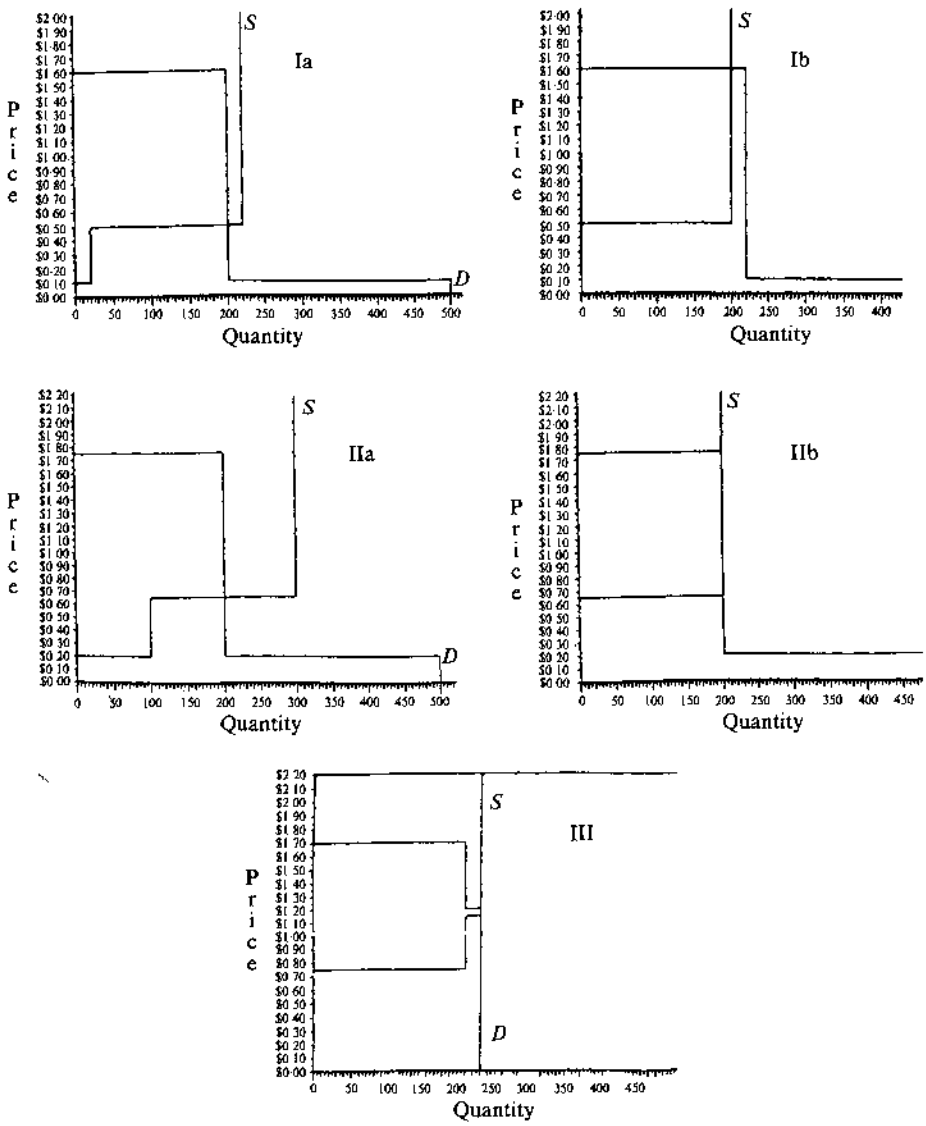

Fig. 6. Parameter schedules.

withholds 20 of 220 shares while all sellers (or all buyers) fully reveal their units, then the efficiency loss is small (less than 1 % of potential gains) but the favourable price change is substantial.

Typically four subjects are assigned to the seller type (cash endowment = 0, share endowment = 200–300) and an equal number to the buyer type (cash endowment = $10,000, share endowment = 0). Hence market supply and demand quantities are typically four times the *per capita* quantities in Fig. 6. Individual subjects' idiosyncrasies may affect observed buyer and seller

© Royal Economic Society 1995

Table 1
Experimental Design Parameters

Session	Experiment	Payoff schedule	Periods (switch)
1.	DA1	Ia	1–4, 13–16
		Ib	5–12 (9)
2.	DA2	Ia	1–4, 13–18
		Ib	5–12 (9)
3.	DA3	IIa	1–2, 9–10
		IIb	3–8 (5), 11–16 (13)
4.	CHS1	Ia	1–4, 13–16
		Ib	5–12 (9)
5.	CH1	Ia + (15¢, 25¢)	1–4, 13–16
		Ib + (15¢, 25¢)	5–12 (9)
6.	CH2	IIa + (5¢, 0¢)	1–2, 9–10
		IIb + (5¢, 0¢)	3–8 (5), 11–16 (13)
7.	CH3a	III	1–10 (6)
	DA4a		11–20 (16)
8.	DA4b	III + (20¢, 20¢)	1–10 (6)
	CH3b		11–20 (16)
9.	CHQ1	Ia + (15¢, 15¢)	1–4, 13–16
		Ib + (15¢, 15¢)	5–12 (9)
	CH4a	Ia + (15¢, 15¢)	17–20
10.	CHQ2	Ia	6–10, 21–30 (21)
		Ib	11–20
	CH4b	Ia	1–5
		Ib	30–39

Note: All experiments use UC Santa Cruz undergraduate students experienced with single-unit versions of the double auction institution. The number of buyers = number of sellers = 4 except in session 10, where it is 3. The payoff schedules are defined in Fig. 5. In the first session, for example, periods 1–4 and 13–16 use the schedule Ia parameters and periods 5–12 use schedule Ib parameters, and buyers' and sellers' roles are switched at the beginning of period 9.

behaviour and profits. To avoid bias and to avoid complaints of unequal opportunities to make money, buyer and seller assignments are simultaneously switched once or twice in all experiments. All subjects had previous paid experience in computerised laboratory markets, mostly in DA asset markets.

Table 1 summarises the overall design. The first three experiments employ the DA odd-lot institution and parameter schedules I and II; the next three used the same parameter schedules (sometimes shifted up or down by 5–25 cents) in CH experiments. There is therefore a natural pairing of these experiments that helps isolate the effects of the market institution from the effects of the parameter schedules. The next two experiments allow us to isolate from the effects of the subject pool as well: both a DA and a CH market were run for 10 periods each in a single session. The last two experiments tested the CHQ variant against CH.

II.4. *Testable Implications*

Differing theoretical perspectives yield sharply different predictions for the exercise of market power in our experiments. The traditionalist view, supported by incomplete information game theory, calls for substantial attempts to

© Royal Economic Society 1995

exercise market power, and consequently non-competitive inefficient outcomes, in economies with parameters v far from V_{*}^{h}. Thus the traditionalist view predicts inefficiency and low trading volume in all trading periods of our experiments except perhaps for periods using schedule IIa parameters. This view suggests no significant effects for the market institution.

The institutionalist view, by contrast, predicts that the experiments using the Double Auction (DA) market institution will produce efficient Walrasian outcomes for any parameter set. It predicts that box (and box-like) parameters will lead to inefficient outcomes for some other trading institutions, including the Clearinghouse (CH) and its variants.

The as-if complete information NE view predicts efficient Walrasian outcomes in all experiments using the CH and CHS institutions (and by extension the DA institution), but predicts inefficient outcomes in CHQ experiments with non-competitive parameters. For symmetry, we introduce a fourth prediction, supported by some literature (e.g. Voltaire, 1759) but little explicit theory: outcomes will be competitive in any organised market, whether it employs the DA or any CH variant and whatever the parameters. We will refer to this last view as 'Panglossian'.

None of our theories has any striking predictions regarding price since most of our parameters have a very wide range of market-clearing (or nearly market-clearing) prices. Therefore our formal analysis will ignore transaction prices, the traditional focus of market experiments.

Because market power is exercised by withholding trade and manifests itself by Pareto-inefficient final allocations, we measure its extent primarily in terms of net trading volume and efficiency. Before conducting experiments the authors disagreed about the likely outcomes but agreed on the benchmark that efficiencies (actual aggregate gains from trade as a percentage of the maximum) and net trade volume less than (in excess of) 75% would be considered evidence for non-competitive (competitive) behaviour. The reader may find the next section more interesting if he or she now pauses for a moment and chooses a benchmark separating Walrasian from non-Walrasian outcomes.

III. RESULTS

We first orient the reader with an informal overview of the experimental results. Next we perform standard statistical tests on trading volume and market efficiency data to evaluate the competing theoretical views. As with most 'formal' tests, economists may wish to take the results with a grain of salt and regard them as more descriptive than inferential. A justification for such scepticism in the present case is the lack of true independence across periods in a given experiment; e.g. uncontrollable group effects could be present. Nevertheless, we believe that standard tests do provide a useful summary of the data. We close the analysis with a statistical examination of individual behaviour in the CH markets to assess the complete information view.

© Royal Economic Society 1995

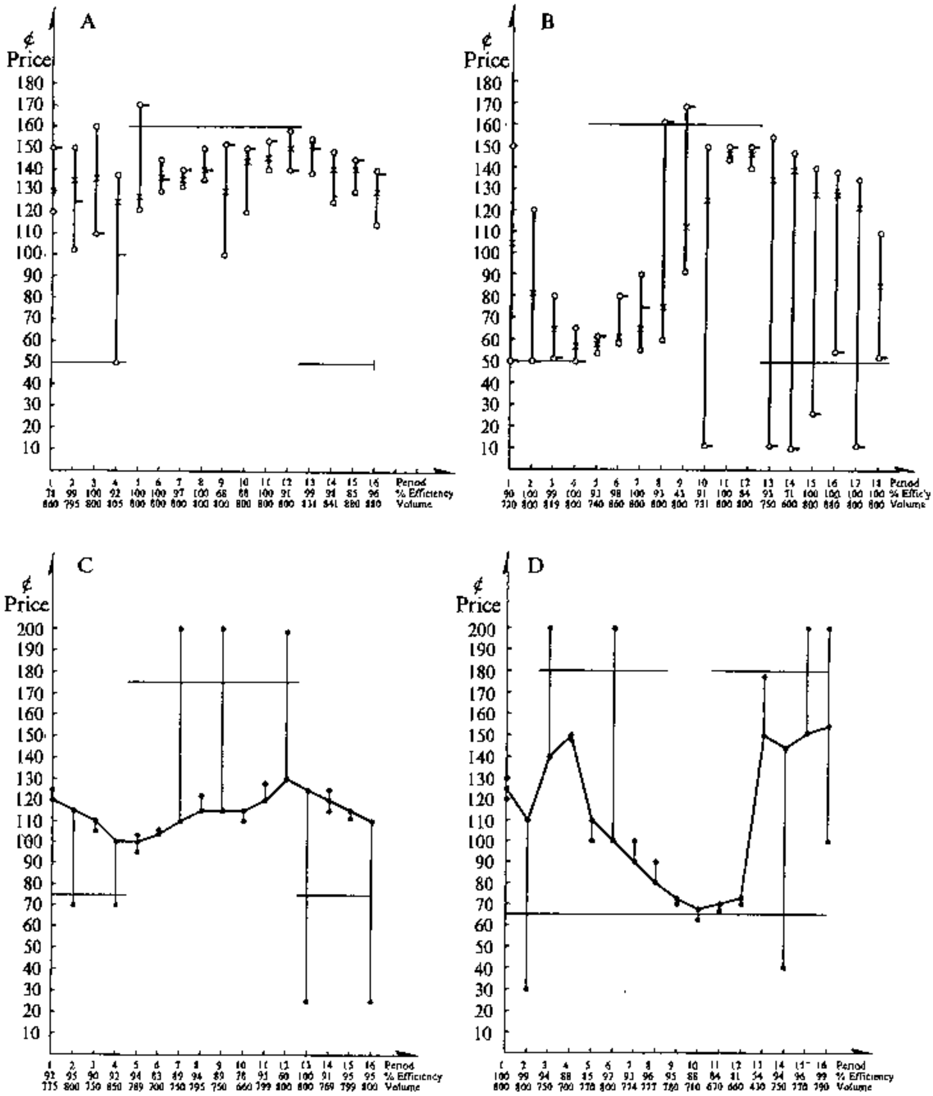

Fig. 7. For legend see facing page.

III.1. *Overview*

Fig. 7 summarises the market outcomes in each trading period of each experiment. For each DA trading period, the endpoints of the vertical line indicate the maximum and minimum transaction prices, the 'x' indicates the mean transaction price and the '–' indicates the close (i.e. final transaction) price. The prices in CH trading periods are similar, except that the clearing prices are connected by straight lines across trading periods, and the best rejected (or nearest extramarginal) bids and asks are indicated by the endpoints of the vertical lines. The longer horizontal lines indicate equilibrium

© Royal Economic Society 1995

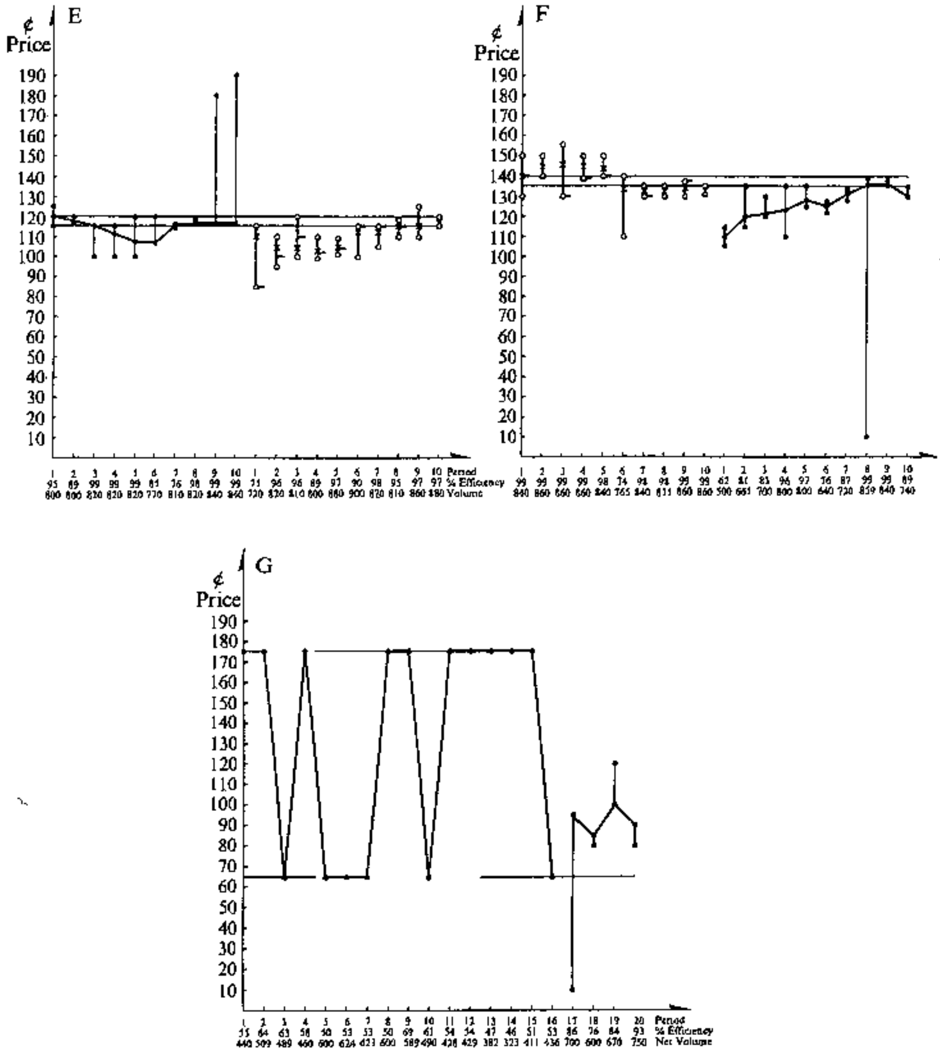

Fig. 7. Experimental outcomes. A, Expt DA1; B, Expt DA2; C, Expt CH1; D, Expt CH2; E, Expts CH3a–DA4a; F, Expts DA4b–CH3b; G, Expts CHQ1–CH4a.

(WE) price. When two horizontal lines extend over a given period, they indicate the maximum and minimum equilibrium price. The two rows of numbers near the bottom of the graph list the percentage efficiency (100%) (actual aggregate trading profit)/(WE aggregate trading profit); and net trading volume, (aggregate absolute difference between initial and final individual holdings)/2, in shares or hundredths of round lots. The Appendix reviews the individual experiments in some detail.

Qualitative impressions of the data may be summarised as follows. Under the DA institution, transaction prices tend to move towards equilibrium (WE) but very slowly, as in previous studies of box-like economies with indivisible units.

© Royal Economic Society 1995

Parameter shifts or buyer/seller assignment switches often increase price volatility and retard convergence. Efficiency is typically above 95 %, with most exceptions due to individual over purchasing errors. There is remarkably little evidence of strategic withholding. Under the basic CH institution the outcomes are roughly similar, with perhaps somewhat lower efficiencies and trading volume, and fewer overpurchasing errors. One sees some possible evidence for strategic withholding with this institution, but it is very sporadic. The two variants of the CH have opposite effects: CHS seems perhaps slightly more efficient and CHQ dramatically less efficient than the basic CH institution.

III.2. *Volume and Efficiency Tests*

Using the 75 % benchmark, the data clearly are inconsistent with the traditionalist prediction, since only 5 of 70 DA trading periods had efficiency at or below 75 %. The data are also inconsistent with a strong version of the institutionalist position, since only 8 of 78 CH (and CHS) trading periods had efficiency at or below 75 %. Moreover, most of the 13 inefficient periods were initial or switch periods. We suspect that few readers will use benchmarks that lead to any different conclusions.

More careful statistical tests will help evaluate the contending theoretical approaches. The standard statistics (see Copeland and Friedman, 1987, for example) include a t-test and Wilcoxon rank-sum test on pooled data (denoted t-pooled and T in Table 2) and a t-test and a signs test on matched pairs

Table 2

Market Performance Comparisons

Comparison							
Pool A (Nobs)	*vs.*	Pool B (Nobs)	Performance measure	t-pooled	T	t-matched	z
1. CHS1		CHQ1	Efficiency	22·64	4·97	20·11	4·00
16		16	Volume	22·29	4·97	19·08	4·00
2. CHS1		DA1	Efficiency	2·05	2·70	1·88	1·50
16		16	Volume	−2·55	−2·56	−2·86	−2·52
3. DA2-4b		CH1-3b	Efficiency	2·06	3·26	2·17	1·73
48		48	Volume	3·11	3·25	3·31	2·59
4. DA4a, b		CH3a, b	Efficiency	1·40	1·09	1·60	0·89
20		20	Volume	2·81	3·02	2·97	2·23
5. CH1, 2, 4		CHQ1-2	Efficiency	10·44	6·87	NA	NA
46		41	Volume	11·75	7·24	NA	NA

Note: Nobs is the number of observations (trading periods) in each pool. The null hypothesis for each statistic is that the efficiency or trading volume measures have the same distribution in Pool A as in Pool B. The statistics are the Student t for equal means (t-pooled), the Wilcoxon rank sum statistic (T), the Student t for differences in matched pairs (t-matched), and the binomial signs statistic for differences in matched pairs (z). The last two statistics are not applicable (NA) when the pools do not come in matched pairs. Positive values indicate a higher mean for Pool A than for Pool B.

© Royal Economic Society 1995

(denoted t-matched and z). A positive (negative) value for any of these statistics indicates that the performance measures for 'Pool A', corresponding to the first level of the treatment variable, are higher (lower) than those for 'Pool B', the alternative treatment. Precise significance levels (given classical assumptions) and asymptotic significance levels can be obtained from standard sources. However, as noted at the beginning of this section, we believe such precision is misplaced and so for descriptive purposes we will refer to absolute values of these statistics in excess of 2·0 (3·0) as significant (very significant).

Table 2 first compares the outcomes in the 16 periods of experiment CHS1 to those in experiment CHQ1. Recall that the two experiments use the same parametric structure but differ in the trading institution in a manner that the as-if complete information approach predicts will be crucial. The results could not have been more clear-cut: every CHS period is more efficient and has higher trading volume than any CHQ period. Hence the non-parametric statistics are as positive as they can be for 16 observations, with the Wilcoxon $T = 4{\cdot}97$ and the signs test $z = 4{\cdot}00$. The t statistics are even more strongly positive, abundantly confirming the dramatic contrast between the extensive exercise of market power in the first CHQ experiment and its virtual absence in the identical-structural CHS experiment.

The next comparison in Table 2 evaluates the institutionalist view by looking for differences between the outcomes in experiment CHS1 and the identically structured double auction experiment DA1. The efficiency test statistics are positive but those for the more reliable paired comparisons are not significant. The negative entries for volume indicate that trading volume is significantly less in the CHS experiment; as explained in the Appendix, the difference is mainly due to a few traders overrevealing quantity (i.e. to obvious blunders) in the DA, and not due to traders strategically underrevealing quantity in the CHS institution.

The next two sets of comparisons evaluate the institutionalist view more generally. Comparing all experiments using the basic CH institution (CH1, 2, 3a, 3b) to the corresponding DA experiments (DA1, 3, 4a, 4b), we obtain positive test statistics generally in the 2·0–3·2 range, indicating significantly greater exercise of market power in the CH institution. A sceptic might object that differing subject pools, and not the market institution, might be the reason for these observed differences; after all, most of the DA experiments were conducted 6 months before the CH experiments and subjects may have become more sophisticated on average.

We address this possibility in sessions 7 and 8 (see Table 1) which use a balanced switchover design to control for subject and group effects. Thus we can be quite confident that the test statistics reported in comparison 4 of Table 2 reflect differences in market institution and not artifacts. Taking into account the smaller number of observations, the results here are very similar to the pooled results in comparison 3.

The final comparison is between all 41 CHQ periods and a pool of 46 of the parametrically most similar CH periods; the unbalanced switchover design of the CHQ experiments precluded matched comparisons. The pooled statistics

© Royal Economic Society 1995

provide highly significant evidence that the CHQ periods have lower efficiency and trading volume. We conclude that all our market performance data support the as-if complete information view over all known rivals.

III.3. *Individual Behaviour in CH Markets*

The strongest evidence in favour of the complete information approach comes from an examination of individual behaviour over time in the CH. Fig. 8 shows

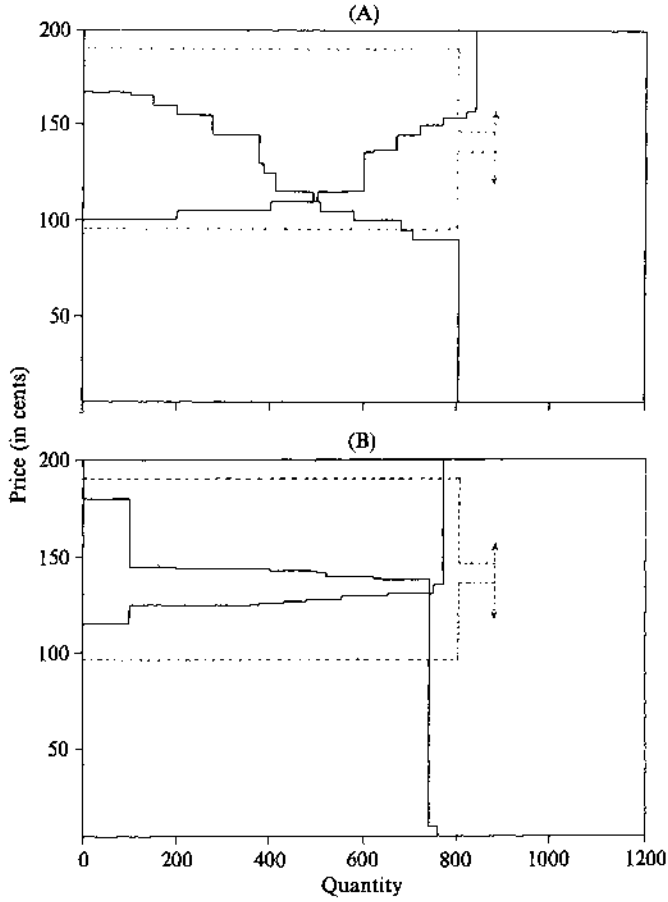

Fig. 8. Supply and demand (A), CH3b, period 1; (B), CH3b, period 10. ....., actual and ---, announced supply/demand.

announced supply and demand for the first and last periods of the last CH experiment, CH3. In general, behaviour in initial periods is quite varied, but that shown in Panel A is not atypical: announced prices (bids and asks) closely approximate true reservation values on some units, are shaded significantly on some other units, and on other units are shaded so much that efficient trades are lost. A recognisable pattern typically emerges after a few trading periods. As in Panel B, most announced (bid and ask) prices are shaded to within 20

© Royal Economic Society 1995

cents of a WE price, although 90% of true reservation values lay 40–50 cents away from the WE price. The result is high efficiency, here about 89%. Perusal of the announced supply/demand diagrams for the other 65 CH market periods (available on request) suggests that, except for the first few periods at the start of an experiment or after a switch of buyer/seller assignments, this pattern is quite typical.

The graphs suggest, then, that traders are converging to the active NE strategies of the DSB theorem, which produce 100% efficiency by shading announced price on inframarginal units all the way to a WE price, i.e. underrevealing price but not quantity. By contrast, in these economies the incomplete information approach suggests a substantial role for underrevealing quantity and perhaps a lesser role for underrevelation of reservation values.

We document price, as opposed to quantity, underrevelation in quantitative terms as follows. Define *RQ* (for revelation of quantity) for a buyer (resp. seller) as the percentage of units valued in excess of (less than) the clearing price which were actually bid (offered). Likewise define *RP* (for revelation of price) for buyers as the mean bid price in excess of the transaction price on units transacted as a percentage of the mean reservation price in excess of the transaction price on those units. More precisely, let

$$qa = \text{no. \{units with reservation price} < pc + 0{\cdot}05\text{\} for sellers}$$
$$(> pc - 0{\cdot}05 \text{ for buyers}),$$

and

$$qr = \text{no. \{units asked at prices} < pc + 0{\cdot}05\text{\} for sellers}$$
$$(> pc - 0{\cdot}05 \text{ for buyers}),$$

where pc is clearing price. It is shaded by 5 cents so units that just missed being transacted are counted, thus tending to produce lower quantity revelation percentages. The latter is defined as $RQ = 100\%\frac{qr}{qa}$. Define price revelation by computing the mean ask (or bid) price, pr, on units transacted (omit the observation if no transactions), and compare to pa, the mean reservation price on these units. Then define price revelation by $RP = 100\%\frac{|pr-pc|}{|pa-pc|}$. Note that the *RP* calculation, unlike the *RQ* calculation, does not involve shading the transaction price; we prefer to err by overstating *RP* and understating *RQ*.

We computed and compared values for *RQ* and *RP* in each trading period for each trader in CH experiments. In the 416 observations, the mean *RP* was 39·6% (standard deviation 30·3), while mean *RQ* was 94·9% (standard deviation 22·9). Table 3 collects some more detailed comparisons. The first line indicates that in 16 of the 16 clearings of CH1, the average value for *RQ* across the 8 traders exceeded the average value for *RP*, with a mean difference of 55·9 percentage points. The null hypothesis that the mean difference is zero is rejected with a t-statistic of 42·89 and a signs test z statistic of 4·00. Similarly overwhelming evidence from other experiments establishes that indeed traders tend to reveal quantity much more fully than price. Overall, mean *RQ* exceeded mean *RP* in each of the 52 relevant clearings, by an average margin of 55·33 percentage points. Indeed, in only 6 of the 202 cases in which they

© Royal Economic Society 1995

Table 3

Price and Quantity Revelation in CH Markets

Sample	Nobs	*RQ–RP*	t	z
CH1	16, 16	55·9%	42·89	4·00
CH2	16, 16	40·8	12·12	4·00
CH3a	10, 10	69·7	18·97	3·16
CH3b	10, 10	63·3	12·97	3·16
Periods 1–8	32, 32	57·1	20·05	5·66
Periods 9–	20, 20	52·5	13·50	4·77
All buyers	202, 194	57·7	20·64	13·09
All sellers	207, 199	53·0	25·15	13·27
All periods	52, 52	55·3	24·04	7·21

Note: The measures of price revelation (*RP*) and quantity revelation (*RQ*) are defined in the text. Nobs is the number of observations (of *RQ*, *RP*); Nobs can differ for the disaggregated data since occasional individual buyers or sellers may be inactive. The *RQ–RP* column reports the difference in sample means for the two revelation measures. The last two columns report values of the matched Student's t and the binomial (z) statistics for the null hypothesis that *RP* and *RQ* have the same distribution.

differed did *RP* exceed *RQ* for individual buyers, and in only 8 of 207 cases for sellers; and even in most of these cases *RQ* was also high. Wilcoxon and Chi-square tests even at the 10% level cannot reject the null hypotheses that *RQ* behaviour is the same for buyers as for sellers, and the same for early as for later periods; and the same statement also holds for *RP*. We conclude that the evidence consistently and firmly supports the view that strategic behaviour in CH markets primarily takes the benign form of price underrevelation rather than inefficient quantity underrevelation, consistent with the complete information approach.

IV. DISCUSSION

We completed two cycles of theory and experiment regarding competitivity and market efficiency with small numbers of traders and odd-lot trading. In the first cycle we addressed the institutionalist versus traditionalist controversy, where the institutionalist position is that competitivity can be enhanced by the form of market organisation and the traditionalist position is that competitivity is largely a matter of configuration of traders' parameter values. Despite our use of non-competitive parameters for induced supply and demand, the market outcomes were highly efficient – and therefore as-if competitive – under the double auction (DA) trading institution (contrary to the traditionalist position) and also under the clearinghouse (CH) institution (contrary to the institutionalist position). On the other hand, the data do support a more cautious statement of the institutionalist view: outcomes are less efficient under the CH institution to a statistically significant (although economically minor) extent.

Examination of individual trading behaviour from the first cycle led us to a second cycle in which we revised our theoretical perspective. We moved away from incomplete information models of non-cooperative equilibrium and adopted as-if complete information as an explanation of the results of the first

© Royal Economic Society 1995

cycle. This also led us to design the CHQ (quantity-only Clearinghouse) market institution which, unlike the CH institution, implies inefficient strategic behaviour at equilibrium. Our two CHQ experiments abundantly confirmed the complete information institutionalist position that the CHQ encourages traders to behave non-competitively and to exercise market power, contrary to the alternative Panglossian position that all organised markets are efficient. These findings raise some fundamental issues which we now address.

IV.1. *On the As-if Complete Information Explanation*

The clearinghouse (CH) data display two striking features. First, the outcomes are highly efficient despite the introduction of odd-lot trade and the choice of non-competitive parameters. Second, supply and demand as revealed in traders' asks and bids typically are highly elastic, so price underrevelation is much more pronounced than quantity underrevelation. These features of the CH data correspond closely to the active NE strategies of the Dubey–Simon–Benassy (DSB) theorem, discussed in Section I.3. In both practice and theory, traders in a CH market prevent others from exercising market power adversely by announcing mutually consistent and highly elastic (competitive) excess demand schedules, even when true excess demand schedules are completely inelastic at the margin.

In Section I.3 we argued that 'complete information' should not be taken literally for theory or for experiment. Our position is that as-if complete information NE provides a plausible explanation of interactions among individuals in certain environments in which information is evidently incomplete. To help make our case, we reconsider previous experimental and theoretical work.

The experimental economics literature contains several precedents for the application of as-if complete information theory to private information experiments, beginning with the classic oligopoly studies of Fouraker and Siegel (1963). The outcomes they report are most consistent with complete information NE in private information experiments, while outcomes in common information experiments, in which all participants' parameters are announced publicly, generally seem to converge toward the cooperative (cartel) solution. Radner and Schotter (1989) provide a more recent example. Although they do not emphasise it in their discussion, outcomes in their bargaining experiments often were more efficient than the incomplete information theory allows. Smith *et al.* (1991) review some older evidence and describe recent oral double auction experiments using indivisible but box-like parameters. They conclude:

> ...contrary to conventional theoretical beliefs, the noncooperative equilibrium concept [complete information NE] fares best in the laboratory under private information [treatments]; under complete (also sometimes common) information [treatments], either the noncooperative equilibrium is not supported (as under two-person bargaining) or its attainment is slower than under private information. [p. 26]

© Royal Economic Society 1995

It appears, then, that there is a pervasive problem in explaining laboratory results. In our view the problem lies in the 'conventional' or orthodox interpretation of games of complete information. In the orthodox interpretation a NE is supported solely by the rational calculations of players who know (as 'common knowledge') the complete structure of the game, including the preferences of all other players. This interpretation is not required by the definition of NE, nor was it always conventional wisdom. An alternative interpretation, that NE arises as a rest point in some groping process, evidently was quite prevalent among game theorists in the 1950s (e.g. Luce and Raiffa, 1957, page 105), but it was largely abandoned after Shapley's negative formal results on fictitious play dynamics reported in Dresher *et al.* (1964). Recently some leading game theorists have come to doubt the orthodox view because it entails logical difficulties related to Gödel's theorem (Binmore, 1987/8) and for other reasons (Kreps, 1990). Binmore's unformalised 'evolutive' processes envisage attainment of complete information NE given very limited information and computational powers (see also the formal processes of Fudenberg and Kreps, 1988, or Friedman, 1991*b*, or in several recent articles in *Games and Economic Behavior*).

Our argument for the as-if complete information model is not based on logical or computational issues, but rather on the fact that in certain dynamic settings, the informational requirements for as-if complete information NE actually can be quite modest. In the present case, traders in a CH market must all anticipate roughly the same clearing price in order to implement strategies of the sort prescribed in the DSB theorem. The laboratory practice of stationary repetition makes such anticipations quite plausible. In most trading periods all traders have the same private values as in the previous trading period, so after experience with stationary repetition they can expect that last period's clearing price is a reliable guide to the current clearing price. Thus, the public announcement of last period's price is really all the information they need to implement the Nash equilibrium strategies prescribed in the DSB theorem. This argument does not apply to the first trading period in an experiment or to trading periods immediately following a shift or switch in private values, but (as noted in Section III) it is precisely in these periods that the complete information theory is least descriptive.

The reader may object that we offer as-if complete information NE as an asymptotic description of market behaviour after it has settled down, but do not offer an explanation of the dynamic equilibration process. We agree, although the fact that efficient trading behaviour is established so rapidly diminishes the importance of this gap. Perhaps new developments in the theory of dynamic or repeated games will someday explain the equilibration process.

The DSB theorem can also help interpret the data from experiments using the other trading institutions. It applies quite directly to the CHS (single offer) variant of the CH institution. The CHQ variant (quantity only) was designed to preclude price-underrevelation strategies of the sort prescribed in the DSB theorem, so we were pleased to see that outcomes under

© Royal Economic Society 1995

the CHQ institution were usually quite inefficient and did not appear to stabilise.[9]

The DA institution is more complex, but the NE strategies prescribed in the DSB theorem may represent some sort of reduced form of DA equilibrium strategies. Certainly traders in the DA can in effect present very elastic excess demand schedules to each other. Indeed the learning required to implement DSB-type strategies could be acquired more rapidly since the information flows (which include unaccepted bid and ask price data as well as timing information and numerous individual transaction price data) are so much richer than in the Clearinghouse.

IV.2. *On the Applicability of the Results*

We have argued that the range of applicability for the incomplete information approach is much broader than generally supposed. The question then becomes, what are its limits? To gather empirical evidence, one could run common information market experiments in which all traders' true parameters are announced publicly. The Smith *et al.* (1991) conclusion cited above suggests that such a move towards more complete information in the experimental procedures might not lead to better agreement of outcomes with the complete information theory. In our view, a more interesting way to probe the limits of applicability is to construct laboratory environments that embody key assumptions of the incomplete information approach. In particular, one could draw parameters randomly each period from publicly announced distributions on a set of V^1 or alternatively V^1_*. Market outcomes from such random values environments should help map the applications boundary between the complete and incomplete information approaches. See Chapters 9 and 10 of Friedman and Rust (1993) for some recent and tentative work in this direction.

Finally, what do the results tell us about everyday markets? Smith (1982*a*) interprets the efficiency of Double Auction laboratory markets as strong support for Hayek's well-known thesis on the informational economy of markets mentioned in the Remark in Section I.3. Our own findings reinforce Smith's by eliminating the alternative interpretation that efficiency in laboratory markets is primarily due to indivisibility and fortuitous choices of parameters. However, the applicability to ongoing markets remains an open question. Some readers (and one author) may be led to believe that most

[9] One might ask why traders do not exercise market power to a greater extent in the CHQ market. After all, efficiencies there are typically 50–75%, rather than 0% that would follow from the box-like parameter configuration. The data show that traders tended to exercise market power, i.e. reduce the quantity demanded or supplied, only when price had become so unfavourable as virtually to eliminate profit. Otherwise, traders tended to increase cautiously the quantity demanded or supplied. What we observed, then, seems to be a two-sided version of 'chiselling' in a cartel. That is, a given buyer finds that the positive effect on prices is the same if a second buyer instead of himself withholds demand and his profit is then greater. Sellers face the same free-rider problem. Hence we should not be surprised that in anonymous market interactions, as in many public goods experiments, we see free-riding; the twist is that here free-riding takes the form of failing to withhold demand or supply.

© Royal Economic Society 1995

contemporary markets can be highly efficient even in the presence of a few large participants if the market is organised by an efficient institution such as the Double Auction.[10] Other readers (and the other author) may disagree, and point out that institutions such as the Double Auction may not even be feasible in important everyday markets for commodities which are not completely standardised. We note that Bertrand-like competition is theoretically possible in such an environment (as in Simon, 1987), but the market institution used in the theoretical demonstration is very complex and perhaps impractical. We agree that it will be important to study laboratory markets with non-homogeneous commodities when advances in laboratory technology make it feasible to conduct such experiments.

University of California, Santa Cruz

University of California, Los Angeles

Date of receipt of final typescript: August 1994

APPENDIX: SUMMARY OF THE EXPERIMENTS

A detailed discussion of individual experiments may be useful since our market institutions and parameters are somewhat unusual. The remarks to follow focus on Fig. 7.

DA1 was a pilot experiment that employed subjects experienced in ordinary asset market experiments (reported in Copeland and Friedman, 1987). These subjects were accustomed to price typically near the highest payoff, and this strong expectation probably accounts for the sizeable upward bias in period 1–4 prices. The reader may detect a tendency for mean transaction prices (marked on the graphs by an '*x*') to converge to equilibrium, but such a tendency is weak at best. Period 9 represents a 'shock': all participants were switched between buyer-type and seller-type value parameters, producing greater price volatility and greatly reduced efficiency. An examination of individual behaviour shows that the cause was gross overpurchasing (i.e. buying a block of shares at a price that initially is profitable but is very unprofitable on the last part of the block) by trader no. 5. Such individual errors (we can think of no rationalisation for such behaviour; subjects' groans at the end of these periods confirmed our belief that the overpurchasing was inadvertent) actually explain all instances of efficiency less than 90% in this experiment. Net trading volume almost always equalled or exceeded the Walrasian equilibrium value of 800 shares.

In practice (no payment) periods before DA2, traders contemplated parameters that would result in equilibrium price of 50¢; perhaps this change in expectations reduced the price bias observed in the previous experiment. A bug in the program allowed short sales on rare occasions;[11] such an occasion presented itself in period 6, increasing the

[10] For example, consider the publicity about manipulation of the primary US Government Securities market by a large participant. If the market were organised as a uniform price auction, such as a clearing-house, rather than its present form (a kind of discriminatory quasi-auction), the incentives and opportunities to manipulate would be considerably reduced. See Chari and Weber (1992).

[11] For the morbidly curious, the bug details are as follows. A single character error in the computer source code allowed the holder of the market ask to sell his last shares by accepting the market bid without automatic cancellation of his ask. Thus sometimes another trader could then accept the old market ask, resulting in a short position for the first trader. This bug was eliminated before any subjects learned how to exploit it.

© Royal Economic Society 1995

supply just enough to make 50¢ an equilibrium price. Except for the resulting delay, the trend toward equilibrium in (mean transacted) price is clear (albeit slow) in this experiment. With the exceptions of period 9 (an assignment switch) and 14 (in which trader no. 4 refused to purchase), efficiency was quite high (89–100%, usually 100% or very close), as was net trading volume.

In experiment DA3, prices clearly trend downward towards the equilibrium value of 60¢ in the IIa periods (1, 2, 9, 10) but otherwise seem to drift in the WE price interval. The lower efficiencies in periods 5 (a switch period) and 6 are again due to overpurchasing errors rather than withholding. In period 15, trader no. 5 did not sell and trader no. 1 did not purchase 30 of the 200 shares predicted in competitive equilibrium; this tiny perturbation is the best evidence of strategic withholding in the whole series of DA experiments.

Recall that the next set of experiments, designated CHS1 and CH1-2, features a single clearing of sealed bids and asks. Only one bid or ask per trader is allowed in CHS1, but up to eight are allowed in CH1 and 2. The efficiencies and volume in CHS1 seem at least as high as those in the corresponding DA experiment, and price also appears to converge, somewhat erratically and very slowly, to equilibrium.

CH1 outcomes are roughly similar, with somewhat lower efficiencies and price convergence perhaps not quite so slow. One observes some apparent withholding, most dramatically in period 6: trader no. 5 did not sell 100 shares, trader no. 0 being the main underpurchaser, and efficiency was only 83%. The depressed efficiency in Period 12 was due to a gross overpurchase error by trader 4. The switch in Period 9 had no obvious impact.

CH2 featured parameter set II with switches in Periods 5 and 13; the impact in this case is evident as traders nos 4–8 were consistently able to get prices to move in their favour. A review of individual bidding behaviour in Periods 8–12 shows that traders nos 0–3 (sellers in these periods) usually revealed supply almost fully with more aggressive asks usually made only by a single seller. Given the elasticity of the rest of the supply curve, these attempts seemed to have little effect other than reducing that seller's meagre profit. The buyers, on the other hand, strongly underrevealed their demand, with traders nos 4 and 5 being particularly aggressive. The lower efficiency in Periods 11 and 12 of 84 and 81% seem a consequence of these buyers trying to push the price below the floor (sellers' Pay1) of 65¢/share. Overall, however, volume and efficiency outcomes in this experiment resemble those in the previous experiment.

The next two experiments employ a 'within-subjects' or 'switchover' design to try to isolate the effect of the market institution. The first of these experiments (CH3a–DA4a) used parameter set III and 10 periods of the DA followed by 10 periods of the CH institution with switches of buyers and sellers in Periods 6 and 16; the other experiment is the same except the DA periods follow the CH periods. One sees little difference between the institutions in terms of price; in most periods, the clearing or average transaction price is within 5¢ of the equilibrium band (itself 5¢ wide). The exceptions are mainly in the earlier periods of the second half experiments (DA4a, CH3b), with prices 10–20¢ lower. However, evidence for lower efficiency and trading volume under the CH institution, although still not dramatic, seems clearer in these experiments.

Finally, CHQ1 was designed to test whether traders would withhold in the CHQ institution, which does not allow traders to underreveal reservation price. Here the results are dramatic: trading volume and efficiency are far lower than in any other design. Perusal of individual bid/ask data suggests a simple rule of thumb that describes most subjects' behaviour: if the clearing price was equal to one's own Pay1 in the last period (so marginal profit was 0 and total trading profit was small) then reduce the bid (or ask) quantity in the current period; if clearing price was favourable (so

© Royal Economic Society 1995

marginal profit was 1·10/share with these parameters) then increase the bid (or ask) quantity. The results, indicated in the next to last panel of Fig. 7, are yo-yo prices and rather trendless low efficiency and volume statistics. As we more fully analysed the data and developed the theory, the empirical contrast with the CH institution became increasingly important. Therefore we recently recreated the computer program (which had been lost) and ran the lengthy experiment CHQ2. It exhibits the same qualitative features, and as can be seen from the last panel of Fig. 7, only 5 of the 25 CHQ periods had efficiency in excess of 75%, versus 12 of the 14 CH periods.

References

Benassy, J. P. (1986). 'On competitive market mechanisms.' *Econometrica*, vol. 54:1, pp. 95–108.

Bertrand, J. (1883). 'Théorie mathématique de la richesse sociale,' *Journal des Savants*, pp. 499–508.

Binmore, K. (1987, 1988). 'Modelling rational players, parts I and II.' *Economics and Philosophy*, vol. 3:2 and 4:1, pp. 179–214, 9–55.

Chari, V. and Weber, R. (1992). 'How the U.S. Treasury should auction its debt.' *Quarterly Review*, vol. 16, no. 4. Minneapolis: Federal Reserve Bank.

Cohen, K. J., Maier, S. F., Schwartz, R. A. and Whitcomb, D. K. (1986). *The Microstructure of Securities Markets*. Englewood Cliffs, NJ: Prentice-Hall.

Copeland, T. E. and Friedman, D. (1987). 'The effect of sequential information arrival on asset prices: an experimental study.' *Journal of Finance*, vol. 42:3, pp. 763–97.

Dresher, M. L., Shapley, L. S. and Tucker, A. W. (1964). *Advances in Game Theory*. Princeton, NJ: Princeton University Press.

Dubey, P. (1982). 'Price-quantity strategic market games.' *Econometrica*, vol. 50:1, pp. 111–26.

Easley. D. and Ledyard, J. (1993). 'Theories of price formation and exchange in double oral auctions.' Chapter 3 in Friedman and Rust (1993).

Fouraker, L. E. and Siegel, S. (1963). *Bargaining Behavior*. New York: McGraw-Hill.

Friedman, D. (1984). 'On the efficiency of experimental double auction markets.' *American Economic Review*, vol. 74:1, pp. 60–72.

Friedman, D. (1991*a*). 'A simple testable model of double auction markets.' *Journal of Economic Behavior and Organization*, vol. 15:1, pp. 47–70.

Friedman, D. (1991*b*). 'Evolutionary games in economics.' *Econometrica*, vol. 59:3, pp. 637–66.

Friedman, D. (1993). 'How trading institutions affect financial market performance: some laboratory evidence.' *Economic Inquiry*, vol. 31:3, pp. 410–35.

Friedman, D. and Rust, J., eds. (1993). 'The double auction market: institutions, theories and evidence.' *Santa Fe Institute Studies in the Sciences of Complexity*, vol. 15. Reading, MA: Addison-Wesley.

Fudenberg, D. and Kreps, D. M. (1988). 'A theory of learning, experimentation, and equilibrium in games.' Mimeo.

Gray, P. and Plott, C. R. (1987). 'The multiple unit double auction.' Caltech Social Science Working Paper no. 625.

Hayek, F. (1945). 'The use of knowledge in society.' *American Economic Review*, vol. 34:4, pp. 519–30.

Holt, C., Langen, L. and Villamil, A. (1986). 'Market power in oral double auctions.' *Economic Inquiry*, vol. 24, pp. 107–24.

Kreps, D. (1990). *A Course in Microeconomic Theory*. Princeton, NJ: Princeton University Press.

Luce, R. D. and Raiffa, H. (1957). *Games and Decisions*. New York: John Wiley & Sons.

Makowski, L. and Ostroy, J. (1987). 'Vickrey–Clarke–Groves mechanisms and perfect competition.' *Journal of Economic Theory*, vol. 42, pp. 244–61.

Makowski, L. and Ostroy, J. M. (1989). 'Efficient and individually rational Bayesian mechanisms only exist on perfectly competitive environments.' UCLA Department of Economics, Working Paper No. 556.

McCabe, K. Rassenti, S. and Smith, V. (1993). 'Designing a uniform-price double auction: an experimental evaluation.' Chapter 11 in Friedman and Rust (1993).

Myerson, R. and Satterthwaite, M. (1983). 'Efficient mechanisms for bilateral trading.' *Journal of Economic Theory*, vol. 29, pp. 265–81.

Plott, C. (1982). 'Industrial organization and experimental economics.' *Journal of Economic Literature*, vol. 20, pp. 1485–527.

Radner, R. and Schotter, A. (1989). 'The sealed-bid mechanism: an experimental study.' *Journal of Economic Theory*, vol. 48, pp. 179–220.

Satterthwaite, M. and Williams, S. (1989). 'Existence and efficiency.' *Journal of Economic Theory*, vol. 48, pp. 107–33.

Satterthwaite, M. and Williams, S. (1993). 'The double auction market: institutions, theories and evidence.' Chapter 4 in Friedman and Rust (1993).

© Royal Economic Society 1995

Simon, L. (1984). 'Bertrand, the Cournot paradigm, and the theory of perfect competition.' *Review of Economic Studies*, vol. 51, pp. 209–30.

Simon, L. (1987). 'Bertrand price competition with differentiated commodities.' *Journal of Economic Theory*, vol. 41, pp. 304–32.

Smith, V. L. (1962). 'An experimental study of competitive market behavior.' *Journal of Political Economy*, vol. 70, pp. 111–37.

Smith, V. L. (1982*a*). 'Markets as economizers of information: experimental examination of the "Hayek hypothesis".' *Economic Inquiry*, vol. 20, pp. 165–79.

Smith, V. L. (1982*b*). 'Microeconomic systems as an experimental science.' *American Economic Review*, vol. 72, pp. 923–55.

Smith, V. L., McCabe, K. A. and Rassenti, S. J. (1991). 'Lakatos and experimental economics.' Chapter 6 in *Appraising Economic Theories: Studies in the Methodology of Research Programs* (ed. N. de Marchi and M. Blaug), pp. 197–226. Aldershot: Edward Elgar.

Smith, V. L. and Williams, A. W. (1989). 'The boundaries of competitive price theory: convergence, expectations, and transactions costs.' In *Advances in Behavioral Economics*, vol. 2 (ed. L. Green and J. H. Kaqel). Norwood, New Jersey: Ablex Publishing Corp.

Smith, V. L., Williams, A. W., Bratton, K. and Vannoni, M. (1982). 'Competitive market institutions: double auctions versus sealed bid-offer auctions.' *American Economic Review*, vol. 72, pp. 58–77.

Van Boening, M. and Wilcox, N. (1994). 'Avoidable cost: ride a double auction roller coaster.' University of Mississippi Department of Economics and Finance, mimeo.

Voltaire, J. F. M. A. de (1928, original published 1759). *Candide.* New York: Random House.

Wilson, R. (1987). 'Equilibrium in bid–ask markets.' In *Arrow and the Ascent of Economic Theory: Essays in Honor of Kenneth J. Arrow* (ed. G. Feiwel). London: Macmillan.

© Royal Economic Society 1995