

The Allocative Effectiveness of Market Protocols Under Intelligent Trading

Marco LiCalzi¹ and Paolo Pellizzari²

¹ Dept. Applied Mathematics and SSAV, U. of Venice, Italy licalzi@unive.it

² Dept. Applied Mathematics and SSAV, U. of Venice, Italy paolop@unive.it

2.1 Introduction

An important criterion for the evaluation of an exchange market is its ability to achieve allocative efficiency. The seminal paper by Gode and Sunder (1993) shows that the protocol known as continuous double auction can attain the efficient allocation even if the traders exhibit “zero-intelligence”: hence, market protocols may actively contribute to the discovery of an efficient allocation. This paper spawned a variety of computer simulations that “enabled us to discover that allocative efficiency [...] is largely independent of variations in individual behavior” at least in canonical environments; see Sunder (2004).

However, the attainment of allocative efficiency is only a necessary condition for the effectiveness of a market protocol in an exchange economy. For instance, consider the fictitious protocol of Walrasian tâtonnement, where a centralized market maker iteratively elicit traders’ excess demand functions and adjust prices before trade takes actually place. Under standard conditions, this protocol attains allocative efficiency while simultaneously minimizing both the volume of transactions and price dispersion. Moreover, the efficient allocation is reached in one giant step, so that its speed of convergence (after trade begins) is instantaneous.

Clearly, the Walrasian mechanism is only an idealization. Realistic market protocols require far less information from traders and should not be expected to perform as smoothly. This raises the question of ranking the effectiveness of those different market protocols which are commonly used in real markets; see Audet et al (2002) or Satterthwaite and Williams (2002). Assuming that they all pass the test of achieving an efficient allocation, which additional criteria should enter in their comparison? Walrasian tâtonnement suggests at least three possibilities: excess volume, time to convergence, and price dispersion.

A major complication in the study of alternative protocols is that their outcome is profoundly affected by traders’ behavior; see Brewer et al (2002). This may exhibit sophisticated strategies, behavioral biases, access to different forecasting abilities, and a variety of factors which we encompass under the

term of traders' *intelligence*. Gode and Sunder (1993) introduced the notion of "zero intelligence" as an extreme assumption, under which all complications in traders' behavior are ruled out and traders are only requested to satisfy a natural budget constraint. They argued that the outcome of a market protocol under zero intelligence is a test of its intrinsic ability to perform effectively.

Assuming zero intelligence, LiCalzi and Pellizzari (2005) compares the performance of different market protocols with regard to allocative efficiency and other criteria such as excess volume or price dispersion. The main protocols examined are: the batch auction, the continuous double auction, a (nondiscretionary) specialist dealership, and a hybrid of these last two. All the four protocols exhibit a remarkable ability to achieve allocative efficiency under three variants of zero intelligence, confirming the main insight from Gode and Sunder (1993).

However, even under zero intelligence, stark differences in performance emerge over other relevant dimensions. The continuous double auction has the worst performance with respect to excess volume, time to convergence, and price dispersion. The dealership has a lower time to convergence and never performs worse than the batch auction. These differences are sometimes dramatic and sometimes small (but persistent). Hence, LiCalzi and Pellizzari (2005) concludes that (under zero intelligence) there is a clear partial ranking of these protocols with respect to excess volume, time to convergence, and price dispersion. A dealership performs slightly better than a batch auction or a hybrid market, and both are substantially more effective than a continuous double auction.

The relevance of this conclusion for the evaluation of practical market protocols is severely limited by the assumption of zero intelligence, which rules out the impact of differences in traders' behavior. The question addressed in this paper is how much of this conclusion remains true if we remove zero intelligence. Using two simple rules for intelligent trading, we study the performance ranking for the four market protocols with regard to excess volume, time to convergence, and price dispersion.

The organization of the paper is the following. Section 2.2 describes the model used in our simulations. Section 2.3 details the experimental design. Section 2.4 reports on the results obtained and Section 2.5 offers our conclusions. For an expanded and more robust analysis, see LiCalzi and Pellizzari (2006).

2.2 The Model

We use the same setup as in LiCalzi and Pellizzari (2005), where a simple exchange economy admits a unique efficient allocation. Given that the market protocols attain allocative efficiency, this implies convergence to the same allocation and facilitates comparisons. Following Smith (1982), we identify

three distinct components for our (simulated) exchange markets. The environment in Section 2.2.1 describes the general characteristics of the economy, including agents' preferences and endowments. The market protocols in Section 2.2.2 provide the institutional details which regulate the functioning of an exchange. The behavioral assumptions in Section 2.2.3 specify how agents make decisions and take actions.

2.2.1 The Environment

We consider an economy with n traders. There is cash and one good, which is an asset with a (random) realization value Y at a given time T in the far future. Each trader i has an initial endowment of cash $c_i \geq 0$ and shares $s_i \geq 0$. We rule out any informational effect and assume that all traders believe that Y is normally distributed with mean $\mu \geq 0$ and precision $\tau = 1/\sigma^2 > 0$ and that no new information is ever released. Therefore, traders' beliefs about Y are homogeneous and never change until uncertainty resolves.

Each trader i has CARA preferences over his final wealth, with a coefficient of risk tolerance $k_i > 0$. Therefore, trader i 's excess demand function for the asset (net of his endowment s_i) is the linear function

$$q_i(p) = k_i\tau(\mu - p) - s_i. \quad (2.1)$$

Let $K = \sum_i k_i$ be the sum of traders' coefficients of risk tolerance. The unique efficient risk-sharing allocation for this economy requires that trader i holds $s_i^* = (S/K)k_i$ asset shares; that is, it is proportional to the coefficient of risk tolerance. The competitive equilibrium achieves the efficient allocation at the price $p^* = \mu - S/(\tau K)$. At such price, the trader i 's net demand $q_i(p^*)$ is exactly filled, making his final allocation $q_i(p^*) + s_i$ equal to the required $s_i^* = (S/K)k_i$.

2.2.2 The Market Protocols

We compare the performances of four market protocols: a batch auction, a continuous double auction, and a nondiscretionary dealership, and a hybrid of these last two. The first protocol is simultaneous, while the other three are sequential. The following features are common to all protocols.

A protocol is organized in trading sessions (or days). Agents participate in every trading session, but each of them can exchange at most one share per session. If the protocol is sequential, the order in which agents place their orders is randomly chosen for each trading session. If the protocol is simultaneous, all order are made known and processed simultaneously so the time of their submission is irrelevant. In every trading session, each agent selects on which side of the market he attempts to place a trade: he can switch roles across trading sessions, but he cannot place simultaneous orders

for buying and selling within the same session. The books are completely cleared at the end of each trading session.

Prices are quoted using a minimum tick; in other words, they are discretized. Moreover, prices must be nonnegative: if a trader places a bid lower than zero, this is ignored; if a trader places an ask lower than zero, this is automatically converted to the lowest strictly positive price compatible with the existing tick.

The specific market protocols studied in this paper are the following.

Batch auction. In each trading session, after traders submit their orders, the exchange price p^* is obtained at the intersection of demand and supply. If there are multiple solutions, we choose p^* as the midpoint of the interval between the lowest and the highest possible values. (If there are no solutions, no exchange takes place.) Shares and corresponding payments are exchanged between traders who submitted bids not lower than p^* and asks not higher than p^* . Traders who placed orders exactly at price p^* may be accordingly rationed. This protocol is also known as the k -double auction, with $k = 1/2$.

Continuous double auction. In each trading session, traders place their orders on the selling and buying books. Their orders are immediately executed if they are marketable; otherwise, they are recorded on the books with the usual price-time priority. Orders are canceled only when a matching order arrives or the trading day is over.

Nondiscretionary dealership. There is a specialist dealer who posts bid and ask quotes valid only for a unit transaction. Agents check sequentially the dealer's quotes for the side of the transaction they are attempting. If an agent accepts the dealer's quote, the exchange takes place at the quoted price. Right after a transaction is completed, the two dealer's quotes for bid and ask increase (or decrease) by one tick if the agent completed a purchase (or a sale). The size of the bid-ask spread stays fixed over time, so the price is never unique. Limited to this protocol, therefore, convergence of prices to a given value p^* should be interpreted as convergence to within a bid-ask interval that contains p^* .

Hybrid market. This combines the continuous double auction with the dealership. Distinct selling and buying books hold quotes from the specialist dealer and from the public, respectively. The dealer posts bids and asks valid only for a unit transaction and revises her quotes as in the nondiscretionary dealership; in particular, she moves her quotes only after transactions in which she has been involved. Agents check sequentially the books for the side of the transaction they are attempting. Their orders are immediately executed at the best price available (which may be different from the specialist's) if they are marketable; otherwise, they are recorded on the traders' book with the usual price-time priority. Agents' orders are canceled only when a matching order arrives or the trading day is over. Hence, once deposited on the traders' book, an order from an agent cannot be executed with the dealer.

2.2.3 Behavioral Assumptions

A major obstacle in the study of microeconomic systems is that their performance is jointly determined by the interactions of traders' behavior within the market protocol. As traders may react differently to different market protocols, it is difficult to separate the intrinsic characteristics of a market protocol from the properties induced by the traders' strategies. Our approach is to concentrate on the institutional characteristics of the protocols, by making general-purpose assumptions on traders' behavior. These assumptions hold for all the simulations reported in this paper.

First, traders are restricted to trade one unit at a time. This restriction on traded quantities simplifies the strategy space and allows direct comparisons with existing theoretical results. Second, buying orders are constrained by the available cash and selling orders by the available endowment of the asset; that is, budget constraints hold. This is consistent with a value-based strategy ("buy low, sell high"), which is a seemingly natural requirement of rationality for traders' behaviors.

Third, since the demand function (2.1) of each trader is strictly decreasing, traders have decreasing marginal utility for additional units. If the current endowment of a trader is s_i , his valuation for the next unit to trade is

$$p(\pm 1) = \mu - \frac{s_i \pm 1}{k_i \tau}, \quad (2.2)$$

where the \pm sign depends on whether the attempted trade is a purchase or a sale. Clearly, this implies that the reservation price of each trader depends on the side of the transaction he is entering and on his current endowment s_i ; moreover, his (implicit) bid-ask spread is $2/(k_i \tau)$.

Given his valuation, in each stage a trader must decide which side of the transaction he wants to attempt and what price to offer. These two separate decisions may exhibit various degrees of intelligence. LiCalzi and Pellizzari (2005) models zero intelligence as follows. At the start of a trading session, each trader chooses either side with equal probability. This randomized choice is stochastically independent of previous history, endowment, or any other parameter of the model. Hence, a trader ignores that the current market price is an imperfect signal for whether he should seek to buy or sell. After the trader has chosen his side of the transaction, suppose that he is going to attempt a purchase. Then his valuation for the next unit to buy is $p(+1)$ from Equation (2.2). Under zero intelligence, this (potential) buyer bids a price uniformly drawn from the interval $[p(+2), p(+1)]$. Similarly, a (potential) seller asks a price uniformly drawn from the interval $[p(-1), p(-2)]$. Again, the information associated with the current market price is ignored.

There are numerous possibilities to make traders "intelligent", ranging from the simple to the highly sophisticated. We attempt to capture the essence of intelligent trading by making two distinct assumptions that exploit the imperfect signal associated with the current price. One is concerned with the

choice of which transaction to attempt, and the other with the price that is offered or asked.

Suppose that the current market price³ is p . An intelligent trader decides the side of his next (potential) transaction by comparing p with his current valuations $p(+1)$ and $p(-1)$. If $p(+1) > p$, the market price is lower than the price at which the trader would like to buy one more unit, so he attempts a purchase. If $p(-1) < p$, the market price is higher than the price at which the trader would like to sell one more unit, so he attempts a sale. If $p(-1) \leq p \leq p(+1)$, the probability that he attempts a purchase is proportional to the distance of p from $p(+1)$. Formally, we assume that with probability

$$P(\text{sale}) = \begin{cases} 0 & \text{if } p(-1) < p \\ \frac{p-p(+1)}{p(-1)-p(+1)} & \text{if } p(-1) \leq p \leq p(+1) \\ 1 & \text{if } p(+1) > p \end{cases}$$

the trader attempts a sale and otherwise goes for a purchase. For later use, we nickname this assumption S_1 as a mnemonic for “side”. The former zero intelligence assumption with $P(\text{sale}) = 1/2$ is denoted S_0 . (The subscripts “1” and “0” denote the presence or absence of intelligence.)

Next, consider the choice of the price. Suppose that the trader is attempting a purchase. Under zero intelligence, he would post a bid uniformly drawn from the interval $[p(+2), p(+1)]$. We model intelligent trading by assuming that he compares the current market price p and his demand function to find the interval $[p(n+1), p(n)]$ which contains p and then posts a bid uniformly drawn from $[p(n+1), p(1)]$. Compared to zero intelligence, this trader selects his bid from a larger and more aggressive interval. There is a nice intuition for this rule: at a price p in $[p(n+1), p(n)]$, the trader would like to buy n units. However, as we constrain him to buy one unit a time, he can at best try to buy the next unit at a price no greater than $p(n+1)$. The symmetric version holds when the trader is attempting a sale. We nickname this rule P_1 as a mnemonic for “price”; P_0 denotes the rule under zero intelligence.

2.3 Experimental Design

2.3.1 Identification

A simulation run for our model requires the specification of five global parameters, a list of individual variables for each trader, as well as specific assumptions about market protocol and traders’ behavior. The global parameters are the number n of traders, the mean μ and the variance σ^2 of the realization value Y of the asset, the number t of trading sessions, and the size Δ of

³ For a batch auction, we use the price observed in the last active trading session. For the sequential protocols, we use the midpoint of the (best available) bid-ask spread.

the tick. Individually, a trader i is characterized by his coefficient k_i of risk tolerance and by his endowment of cash c_i and asset shares s_i . Finally, for protocols involving the dealer, we need to select her initial quotes.

The market protocols are described in Section 2.2.2. For ease of reference, we nickname the protocols as B (batch auction), C (continuous double auction), D (automated dealership), and H (hybrid market). Recall that the behavioral assumptions described in Section 2.2.3 are nicknamed S_i and P_k for $i, k = 0, 1$.

We have run simulations for all $4 \times 4 \times 3 = 48$ possible combinations of protocols, behavioral assumptions and performance criteria, over different instantiations of the parameters. The results reported in Section 2.4 are robust both to variations in the fine details in the market protocols and substantial changes in the parameters, provided that the overall liquidity of the system is sufficiently large. To simplify the presentation, we fix the exemplar parametric configuration reported in Table 2.1 and for each performance criterion we report the simulations for the four market protocols and the four behavioral assumptions. The initial dealer's quotes are a bid of 745 and an ask of 751, with a fixed bid-ask spread of 6. The competitive equilibrium price is $p^* = \mu - 2\sigma^2 = 760$ in all the simulations reported in this paper.

Table 2.1. Exemplar for identification.

	Parameters	Initialization
Global	n	= 1,000
	μ	= 1,000
	σ^2	= 120
	t	= 2,500
	Δ	= 1
Trader	k_i	= divisors of σ^2 in $\{10, \dots, 40\}$
	c_i	= 50,000
	s_i	= permutation of $2k_i$

We say that a market protocol *exactly implements* a trading rule if it is never necessary to round traders' offers to match the ticked prices; see LiCalzi and Pellizzari (2005). An exact implementation allows *exact* convergence to the equilibrium price supporting the efficient allocation. (This is not relevant for the dealership protocol, because the fixed bid-ask spread prevents the price from being unique.) Our exemplar case is chosen to ensure that all the simulations reported in this paper satisfy the requirement of exact implementation. To this purpose, we choose integer values for μ and σ^2 and initialize each k_i 's with a stochastically independent draw from a uniform distribution over the divisors of σ^2 . We also assume that the support of the uniform distribution over bids and asks is formed by the integers in the two intervals.

2.3.2 The Simulations

A round of testing requires to simulate $3 \times 4 = 12$ combinations of performance criteria and behavioral assumptions. A typical round of simulations runs as follows. For each of the 12 combinations, we instantiate parameters according to the exemplar in Table 2.1 and work out a simulation batch consisting of 100 runs under different initial random seeds. Then we record the time series for prices, volume, and endowments, and compute relevant statistics for the performance criteria. The simulations have been run using a dedicated package of routines written in Pascal.

2.4 Results

We separately evaluate the performance of the four market protocols with respect to three criteria: excess volume, time to convergence, and price dispersion. Each one is defined and discussed in one of the following three subsections.

2.4.1 Excess Volume

Getting from the initial endowment to the efficient allocation requires a minimum number of (unit) transactions. The traded volume is the total number of unit transactions completed before attaining the efficient allocation. The Walrasian protocol attains the efficient allocation in one step and thus minimizes the traded volume. Realistic market protocols usually waste transactions and thus require higher volumes. We measure the *excess volume* in a market protocol as the percentage of traded volume in excess of the minimum required to attain the efficient allocation. Clearly, higher excess volumes signal less effective protocols that let unnecessary trades take place.⁴

Figure 2.1 shows four boxes. Each box is associated with a different assumption about the intelligence of the traders, as noted at its bottom. For instance, the top-right box is associated with S_1P_0 : this corresponds to positive intelligence in the choice of the side and zero intelligence in the price decision. Within each box, we graph the excess volumes for 100 runs for each of the four protocols, as well as marking the average level. The dots are color-coded: black is Batch (B), red is Continuous Double Auction (C), green is Dealership (D), and Blue is the Hybrid (H) protocol. The market protocols perform quite differently and these differences persist under various forms of trading intelligence.

Figure 2.2 merges the dots from the four boxes of Figure 2.1 in a single box. Two main findings emerge. First, the batch auction and the dealership

⁴ When a dealer mediates the transfer of one unit from a trader to another one, we record only one transaction so that the statistics for excess volume are directly comparable.




Fig. 2.1. Excess volume — datapoints.

have a substantially lower excess volume than the other two protocols under any of our variants of intelligent trading; regardless of these, the former two never exhibit more than 4% excess volume, while the latter two never go below 40%. Second, increasing trading intelligence tends to reduce excess volume, most notably in the continuous double auction, but does not eradicate the differences.

We fit a linear model to the data using a robust regression based on an M estimator; see Venables and Ripley (2002). The independent dummy variables are B, C, D , and H for the protocols, and P, S for trading intelligence over price and side. Dummies for protocols are increasingly ordered by the size of their effect: here, we leave out D because it has the lowest marginal impact. With t -values reported below each coefficient, the estimated equation for the excess volume is

$$\text{ExcVol} = 0.0566 + 0.0065B + 0.7192C + 0.8205H - 0.0252S - 0.0889P$$

(16.70) (1.673) (183.9) (209.8) (-9.126) (-32.13)

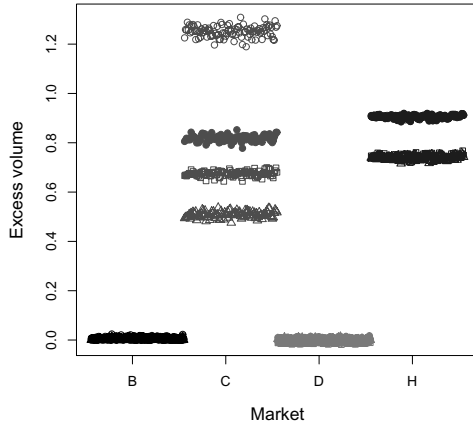


Fig. 2.2. Excess volume — merged datapoints.

The intercept, of course, combines the joint effect of D and S_0P_0 and thus the baseline is a dealership protocol with zero intelligence trading. It is clear that the (average) effect of C and H on increasing the excess volume is statistically significant. Similarly, side and price intelligence decrease the excess volume.

2.4.2 Time to Convergence

Our second performance criterion is the number of trading sessions completed before no further trading takes place. In our exemplar case, the maximum number of units between the initial endowment and the final efficient endowment is 60, so this is a lower bound on the number of trading sessions required to achieve allocative efficiency. Figure 2.3 is similar to Figure 2.2 and reports the merged datapoints for time to convergence.

The estimated equation for time to convergence is

$$\text{Time} = 183.31 + 16.022H + 156.78C + 208.98B - 149.26S + 7.5212P$$

$$(45.99) \quad (3.481) \quad (34.06) \quad (45.40) \quad (-45.86) \quad (2.311)$$

The (average) effect of H, C and B on increasing the time to convergence is statistically significant. Remarkably, while side intelligence contributes to this reduction, the coefficient for price intelligence denotes a (weak) contrary effect — when trading is aggressive, time to convergence lengthens.

2.4.3 Price Dispersion

Our third and final performance criterion attempts to quantify the dispersion of prices by measuring the standard deviation of the time series of the prices




Fig. 2.3. Time to convergence — merged datapoints.

observed at the end of each trading session. Figure 2.4 reports the merged datapoints for price dispersion.




Fig. 2.4. Price dispersion — merged datapoints.

The estimated equation for price dispersion is

$$\begin{aligned}
 PDisp = & 1.5128 + 1.3868H + 3.0874D + 15.398C + 0.2826S - 0.5027P \\
 & (29.61) \quad (23.50) \quad (52.32) \quad (261.0) \quad (6.773) \quad (-12.05)
 \end{aligned}$$

The (average) effect of H , D and C on increasing price dispersion is statistically significant, and particularly sizeable for C . Side intelligence has also an increasing effect, while price intelligence has a moderating impact.

2.5 Conclusions

The simulations shows that the choice of a protocol may have a substantial impact on the allocative effectiveness of an exchange market. Since lack of space prevents us from a longer analysis, we offer only the main conclusions. A richer and more complete analysis is carried out in LiCalzi and Pellizzari (2006).

Excess volume. The ranking with respect to excess volume is $\{B, D\} \gg C > H$, where $>$ stands for “lower volume” and \gg for “much lower volume”. The notation $\{B, D\}$ means that the ranking is not statistically significant. In simple words, the batch auction and the dealership generate minimal excess volume; on the other hand, protocols involving a continuous double auction are seriously wasteful. Moreover, intelligent trading helps, in the sense that increasing the intelligence of traders tends to reduce (but does not eradicate) the excess volume.

Prescriptively, this suggests that a market regulator attempting to reduce excess volume in an exchange market would be well advised to opt for a batch auction or a dealership. Moreover, he should make an effort to educate traders towards making use of the signals embedded in the market price.

Time to convergence. The ranking with respect to time to convergence is $D > H \gg C > B$, where $>$ stands for “lower time” and \gg for “much lower time”. Protocols involving a dealer converge much faster. Intelligent trading is overall beneficial but has an ambiguous effect. A better choice for the side of the transaction to attempt substantially reduces the time to convergence: this alone might wipe out differences among all protocols except for the batch auction. On the other hand, more aggressive behavior on the choice of the prices slightly increases this time.

Prescriptively, this suggests that a market regulator attempting to reduce the time to convergence in an exchange market should consider having a dealership (possibly along an open book). Moreover, he should point out to traders the importance of using the price signal to understand the direction in which trade should be oriented, while attempting to reduce their greediness.

Price dispersion. The ranking with respect to price dispersion is $B > H > D \gg C$, where $>$ stands for “lower dispersion” and \gg for “much lower dispersion”. The batch auction minimizes price dispersion and the continuous double auction yields by far the worst performance in this respect. Intelligent trading is overall damaging but with an ambiguous effect. More intelligence on choice of the side of the transaction increases the dispersion, while a more

aggressive pricing behavior has a mild moderating effect. Prescriptively, this suggests that a market regulator attempting to reduce price dispersion in an exchange market should avoid the use of a continuous double auction.

References

- [1] Audet N, Gravelle T, Yang J (2002) Alternative trading systems: Does one shoe fit all?. Working paper 2002-33, Bank of Canada, November
- [2] Brewer PJ, Huang M, Nelson B, Plott CR (2002) On the behavioral foundations of the law of supply and demand: Human convergence and robot randomness. *Experimental Economics* 5:179-208
- [3] Gode DK, Sunder S (1993) Allocative efficiency of markets with zero intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy* 101:119–137
- [4] LiCalzi M, Pellizzari P (2005) Simple market protocols for efficient risk sharing. RePEc:wpa:wuwpfi:0504019, April.
- [5] LiCalzi M, Pellizzari P (2006) The allocative effectiveness of simple market protocols. Working paper.
- [6] Satterthwaite MA, Williams SR (2002) The optimality of a simple market mechanism. *Econometrica* 70:1841–1863
- [7] V. L. Smith (1982), “Microeconomic systems as an experimental science”, *American Economic Review* **72**, 923–955.
- [8] Sunder S (2004) Markets as artifacts: Aggregate efficiency from zero intelligence traders. In: Augier ME, March JG (eds), *Models of Man: Essays in Memory of Herbert A. Simon*. The MIT Press, 501–520
- [9] Venables WN, Ripley BD (2002) *Modern Applied Statistics with S*. Fourth edition. Springer.