6

Which Market Protocols Facilitate Fair Trading?

Marco LiCalzi and Paolo Pellizzari

Dept. Applied Mathematics and SSAV, U. of Venice, Italy
[licalzi,paolop]@unive.it

6.1 Introduction

The evaluation of an exchange market is a multi-faceted problem. An important criterion is the ability to achieve allocative efficiency. Gode and Sunder (1993) shows that a continuous double auction for single-unit trades leads to an efficient allocation even when the traders exhibit "zero-intelligence"; in other words, market protocols are active contributors in the search for a better outcome. Under reasonable circumstances, most of the commonly used market protocols share the ability to help traders discover an efficient allocation.

As suggested in Hurwicz (1994), however, the attainment of allocative efficiency is only a necessary condition for the effectiveness of a trading protocol and one should take into account other dimensions. Assuming zero intelligence, LiCalzi and Pellizzari (2007) compares the performance of different market protocols with regard both to allocative efficiency and other criteria such as excess volume or price dispersion. Their study considers agents with decreasing marginal utility that can repeatedly make single-unit trades and examines four common protocols: batch auction, continuous double auction, nondiscretionary dealership, and a hybrid of these latter two. All protocols exhibit a remarkable capacity to achieve allocative efficiency. However, stark differences in performance emerge over the other dimensions. These differences persist even when the assumption of zero intelligence is removed; LiCalzi and Pellizzari (see 2006).

The general conclusion is that although common market protocols may be close substitutes in helping (even zero-intelligent) traders to attain efficiency, they behave quite differently in many respects. This paper expands this line of research moving from the evaluation of al-

locative effectiveness to the assessment of allocative fairness. See Fehr et al. (1993) for a different line of attack on this theme.

Any trading protocol that attains allocative efficiency has two effects. From a static point of view, it moves the traders from their initial endowment to a final (efficient) position where no further paretian improvements are possible and all gains from trade are realized. This ability to help traders discover and exploit all gains from trade pertains to the allocative effectiveness of a market protocol. From a dynamic point of view, on the other hand, the denouement of a trading session decides how these gains are distributed among the traders. The performance of a trading protocol in this respect pertains to its allocative fairness.

A protocol that is allocatively efficient never leaves unrealized gains from trade. A protocol that is allocatively fair makes sure that these gains are equitably distributed among the traders. While many definitions of equitability are possible, there is a general sense that each traders should be entitled to a share of the gains from trade that his being in the market creates. In this paper, we consider the same four common protocols that we have shown to be allocatively efficient (even under zero intelligence) and we ask the following question. Suppose that the market is populated with only two families of agents. Both families are using trading strategies that are individually rational, but the second family enjoys a potential trading advantage on the first one. Which of these market protocols is more effective in making sure that the first family of agents overall loses the least on his "fair share" of gains from trade?

The organization of the paper is the following. Section 6.2 describes the model tested in our computational experiments and formalizes our research question. Section 6.3 details the experimental design and provides detailed instructions for its replication. Section 6.4 reports on the results obtained and Section 6.5 offers our conclusions.

6.2 The model

We use the same setup as in LiCalzi and Pellizzari (2007), where a simple exchange economy admits a unique efficient allocation for the single good to be traded. Given that the market protocols attain allocative efficiency, this implies convergence to the same final allocation of the good and facilitates comparisons.

6.2.1 The environment

We consider an economy with n traders. There is cash and one good, which we call "stock". Each trader i has an initial endowment of cash $c_i \geq 0$ and shares $s_i \geq 0$. 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 stock (net of his endowment s_i) is the linear function

$$q_i(p) = \tau k_i(\mu - p) - s_i. \tag{6.1}$$

where μ is the mean and $\tau = 1/\sigma^2$ is the reciprocal of the variance (a.k.a. as the "precision") of the distribution of the final value of the stock. Each trader knows μ and τ as well as his endowment and his demand function, but otherwise has no information on the other agents.

Let $K = \sum_i k_i$ be the sum of traders' coefficients of risk tolerance, while $S = \sum_i s_i$ and $C = \sum_i c_i$ are the total stock and cash endowments. The unique efficient allocation of shares in this economy requires that trader i holds $s_i^* = (S/K)k_i$ shares of the stock. This is also achieved in the (unique) competitive equilibrium at price $p^* = \mu - S/(\tau K)$; see Wilson (1968). Clearly, the unique efficient allocation of shares is associated with a continuum of feasible allocations for cash; each of these determines a different apportionment of the gains from trade. Therefore, allocative efficiency corresponds to handing out stock in a unique way; allocative fairness has to do with how cash is redistributed during the trading that takes place before the efficient stock allocation is attained.

We emphasize that our setup is not meant to replicate the structure of a stock market; in particular, informational effects are ruled out. The underlying economy can be described as an exchange market for one good, where traders have strictly decreasing linear demands and heterogeneous preferences that are driven by a particularly simple parameterizations.

6.2.2 The market protocols

We compare the performances of four market protocols: a batch auction, a continuous double auction, 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. See LiCalzi and Pellizzari (2006, 2007) for a complete description of the protocols and details on their implementation.

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. Reaching an efficient allocation requires multiple rounds of trading. 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. The books are completely cleared at the end of each trading session. Prices are ticked and, for convenience, the tick is set equal to 1; in other words, prices must be integers.

6.2.3 Behavioral assumptions and fair shares

The following behavioral assumptions hold for each trader. An agent is restricted to trade one unit at a time. Budget constraints must be satisfied. Given the demand function (6.1), trader i has decreasing marginal utility for additional units. If the current endowment of a trader is s, his valuation for the next unit to trade is

$$v_i(\pm 1) = \mu - \frac{s \pm 1}{\tau k_i} \qquad (6.2)$$

where the $\pm$ sign depends on whether the attempted trade is a purchase or a sale. Hence, his reservation price depends on the side of the transaction he is entering and on his current endowment s_i. Moreover, his certainty equivalent for holding quantities c and s of cash and stock is

$$m_i(c, s) = c + \left(\mu - \frac{s}{2\tau k_i} \right) s \qquad (6.3)$$

It is worth noting that the certainty equivalent m_i accounts for c at face value but evaluates s using an individual "price of risk" $\mu - [s/(2\tau k_i)]$.

The initial endowment (c_i^0, s_i^0) of a trader i provides him with a certainty equivalent $m_i^0 = m_i(c_i^0, s_i^0)$. We define his "fair share" m_i^* of gains from trade as the certainty equivalent he would attain under the fictitious protocol of Walrasian tâtonnement, where a centralized market maker iteratively elicit traders' excess demand functions and keeps adjusting prices to equilibrate them *before* trade actually take place. Under standard conditions, this protocol is a natural benchmark because it attains allocative efficiency in one giant step, while simultaneously minimizing both the volume of transactions and price dispersion. For later use please note that, under this protocol, a trader with *ex ante* knowledge of the equilibrium price p^* would attain exactly the

same final certainty equivalent and thus would not be able to increase his fair share.

Under the Walrasian protocol, a trader i ends up with cash $c_i^* = c_i^0 - p_i^*(s_i^* - s_i^0)$ and stock $s_i^* = (S/K)k_i$. After substitution, the certainty equivalent of his fair share is

$$m_i^* = c_i^0 - p^*(s_i^* - s_i^0) + s_i^*\left(\mu - \frac{s_i^*}{2\tau k_i}\right) =$$
$$c_i^0 + \left(\mu - \frac{S}{2\tau K}\right)s_i^0 + \frac{S}{2\tau K}\left(\frac{S}{K}k_i - s_i^0\right)$$

which nicely decomposes into the sum of three terms. The first one is the initial cash endowment of trader i; the second is the "value" of his initial stock endowment at the market price of risk; the third one is a positive correction term that is increasing in the difference between the efficient and the initial stock endowment for i. Since trading is voluntary, individual rationality implies that the difference between the fair share and the initial certainty equivalent for each i is positive:

$$m_i^* - m_i^0 = \frac{\left(Ks_i^0 - k_iS\right)^2}{2\tau k_i K^2} \geq 0$$

We expect that market protocols affect how much of their fair share different families of agents manage to obtain in the end. This requires to aggregate social welfare over groups of agents. We measure the social welfare of a group G by the sum of the certainty equivalents across the traders in G. Given the initial endowments (c_i^0, s_i^0) of each trader i, the (initial) social welfare of the entire traders' population is $M^0 = \sum_i m_i(c_i^0, s_i^0)$. After reaching an efficient allocation, the social welfare increases to

$$M^* = \sum_i m_i^* = C + \left(\mu - \frac{S}{2\tau K}\right)S \qquad (6.4)$$

which is the analog of Equation (6.3) at the market level. We slightly abuse notation here, because M^* is achieved by any efficient allocation including (but not limited to) the one induced by the Walrasian procedure. Looking at the left-hand side of Figure 6.1, efficient trading expands the pie from M^0 to M^*.

Consider now a strict subset G of traders. They start with an initial endowment that corresponds to a social welfare $M_G^0 = \sum_{i \in G} m_i^0$ for the group G. The fair share of this group is $M_G^* = \sum_{i \in G} m_i^* \geq M_G^0$. In the right-hand side of Figure 6.1, we represent M_G^0 as the circular sector from the inside circle and M_G^* as the union of M_G^0 and the annular sector

Fig. 6.1. Gains from trade and fairness.

topping it. In general, M_G^* expands but need not be proportional to M_G^0. Suppose now that at the end of a trading protocol, the social welfare of a group G is $M_G^* \cup A_G$ so that the group G is extracting higher gains from trade than its fair share. Then we say that the protocol has been too favorable to the traders in G or, equivalently, that it has been unfair to the traders in the complementary set G^c. Hence, allocative fairness is about how the larger pie created by trading is redistributed among different groups of traders. Similarly to a zero-sum game, a trader gets more than his fair share by taking away a piece of someone's else fair share.

Our approach to study allocative fairness is to split the traders' population into two families and compare the ability of market protocols to prevent one group from exploiting the other one. For realism, we assume that all agents are individually rational: regardless of which family he belongs to, each agent accepts a trade only if this cannot decrease his current certainty equivalent. An agent who undertakes a sequence of trades over time increases (possibly, weakly) his own certainty equivalent in each transaction. This assumption, for instance, is consistent with zero-intelligence.

Our two families of interacting traders are chosen to emphasize differences in the ability to appropriate gains from trade. Notably, individual rationality alone cannot prevent a purchase from an inframarginal seller even if this reduces the potential gains from a specific trade. Put differently, individual rationality protects a buyer from making a personal loss on a trade but does not imply that he is trading with the "right" counterpart. This stronger guarantee requires knowledge of the equilibrium price p^* in order to spot and refuse inframarginal trades. We assume that some traders satisfy only individual rationality while others can do better because they know p^* as well.[1]

The first group is formed by the truth-telling (from now on, TT) traders described in LiCalzi and Pellizzari (2007). At the start of a

---

[1] An alternative assumption is that only the second type of traders are able to compute or deduce p^* from the available information.

trading session, a TT trader chooses with equal probability on which side of the market (buy or sell) he attempts to trade one unit. Suppose he goes for a purchase; the case of a sale is analogous. Given his current endowment, the agent knows that his valuation for the next unit to buy is $v_i(+1)$ from Equation (6.2). In a batch auction, he truthfully bids $v_i(+1)$. In a sequential protocol, he checks first if the best current ask price is $p \leq v_i(+1)$; if so, he buys one unit at p. Otherwise, he places a bid equal to $v_i(+1)$. In other words, when no better deal is available, a TT buyer posts a bid equal to his current valuation for the next unit to buy and thus "truthfully" reveals his reservation value. Compared to zero-intelligence trading, a TT agent is less greedy because he posts the largest bid that is individually rational given his own valuation. When a TT agent buys one unit at a price p higher than the equilibrium price p^*, he increases his certainty equivalent but eats up a piece $(p - p^*)$ of his fair share.

The second group of agents consists of traders that know the correct equilibrium price; we call them price-informed (from now on, PI). This extreme assumption is a very parsimonious way to endow these agents with the ability to cut down on inframarginal trades and make sure that they never lose on their fair share. Given his current endowment, a PI agent knows that he should be a buyer if $v_i(+1) \geq p^*$ and a seller if $v_i(-1) \leq p^*$. Therefore, he never needs to guess which side he should take.

Suppose that the PI agent should be a buyer; the opposite case is analogous. In a batch auction, he simply bids p^*. In a sequential protocol, a PI trader must take action when he is called out and cannot wait for better terms. When it is his turn, he first looks for "sure deals" by checking whether the best current ask price is $p_a \leq p^*$ or the best bid price is $p_b \geq p^*$; if so, he buys or sell one unit, respectively. Otherwise, and limitedly to the two book-based protocols, a PI agent places a bid that improves the current best bid p_b by one tick and achieves time-price priority at a buying price never greater than p^*. In general, the trading strategy of a PI agent has three characteristics: first, he never fails to exploit opportunities for trading off the equilibrium price; second, he never trades at a price worse than p^* (and hence never loses on his fair share); third, conditional on these two constraints, he maximizes the probability of trading in the right direction. This last restriction is chosen to emphasize the ability of PI traders to take advantage of TT agents.

Depending on the protocols and the random sequence of trades, the attainment of full allocative efficiency may sometimes fail. For

instance, in the nondiscretionary dealership, the existence of a fixed bid-ask spread may prevent two or more TT agents from completing their few last trades. This may (albeit marginally) reduce the overall gains from trade and lower allocative efficiency, confusing our study of allocative fairness. To rule out this spurious effect, after all trading opportunities within the protocol are exhausted, we force agents to carry out all residual efficient trades at price p^*. We emphasize that this has the only purpose of actually realizing the full pie M^* so that we can concentrate on its redistribution; in particular, none of these final trades eats up on the fair share of a trader.

Let M_G^T be the final fair share of a group G when trading takes place using a trading protocol T. Given their information and trading strategies, only PI agents can "exploit" TT traders. Therefore, whenever allocative efficiency is attained, $M_G^T \geq M_G^*$ for $G = PI$ and any protocol T among the four we consider. We can thus test the ability of a trading protocol T to foster a fair allocation by comparing $M_G^T - M_G^*$ for $G = PI$.

Clearly, the ability of the PI group to exploit TT traders depends also on the proportion π of PI traders in the market. The more the exploiters, the harder becomes the competition for trades at prices different from p^*. Therefore, we study how allocative fairness is affected by the proportion π in $(0,1)$. Endpoints of the interval are ruled out to avoid trivialities.

6.3 Experimental design

6.3.1 Identification

The global parameters are the number n of traders, the mean μ and the variance σ^2 of the realization value of the asset, the number t of trading sessions, and the number λ of PI traders. (The proportion of PI agents is $\pi = \lambda/n$.) 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 and a (fixed) spread.

The exemplar for our simulations is similar to that one used in LiCalzi and Pellizzari (2006). The basic parametric configuration is reported in Table 6.1. The ratio $S/K = 2$ implies that the competitive equilibrium price is $p^* = \mu - \sigma^2(S/K) = 760$. The initial dealer's quotes in the nondiscretionary dealership are a bid of 755 and an ask of 765, with a fixed bid-ask spread of 10. In the hybrid protocol, where

	Parameters		Initialization
Global	n	$=$	$1,000$
	μ	$=$	$1,000$
	σ^2	$=$	120
	t	$=$	500
	λ	$=$	integer in $(0, n)$
Trader	k_i	$=$	divisors of σ^2 in $\{10, \ldots, 40\}$
	c_i	$=$	$50,000$
	s_i	$=$	permutation of $2k_i$

Table 6.1. Exemplar for identification.

the dealer's presence restricts the ability of PI traders to steal better deals, the initial bid and ask prices of 745 and 775 exactly straddle the equilibrium price of 760, with a fixed spread of 30.

The robustness tests reported in Section 6.4.1 change one parameter at a time with respect to this exemplar. We have worked out simulations where the ratio S/K is 1 (or 3), making the equilibrium price higher (lower); where the dealer's fixed spread in the nondiscretionary dealership is 6 (or 30), making the market more (less) liquid; and where the fixed spread in the hybrid protocol takes different values between 4 and 300, making the dealer's presence more or less influential.

6.3.2 Simulations and data representation

A round of testing simulates traders' behavior in 4 different protocols for different values of λ. A typical cycle is run as follows. We fix an integer value of λ in the range $\{1, \ldots, n-1\}$ and then we randomly choose different queues of traders for each trading session. These choices are kept fixed across the four protocols, so that each of them is tested using the same fraction of PI traders and the same orderings in each trading sessions. All other parameters are instantiated as per the exemplar in Table 6.1. The number of agents is $n = 1000$; we run 999 trials per cycle and test each value of λ from 1 to 999. At the end of each simulation, we compute and record all relevant statistics. The simulations are run using a package of routines written in Pascal. The statistical and graphical analysis of the data are made using R, an open-source environment for statistical computing available at `http://www.r-project.org/`.

We use two (normalized) measures to assess the allocative fairness of a protocol. Let M_G^T be the final share of the group G when trading takes place using a trading protocol T and M_G^* their fair share (using

the Walrasian protocol). As discussed above, only PI agents can "exploit" TT traders; hence, we fix $G = PI$ for the rest of the paper. The first measure is the absolute *excess gain* $(M_G^T - M_G^*)/M^0$ for the group G. The division by the size M^0 of the initial pie is a normalization introduced to make the index scale–free and allow direct comparisons; however, for simplicity, in the rest of the paper we write the absolute excess gain as $M_G^T - M_G^*$ and leave the normalization implicit. The second measure is the relative *excess gain* $(M_G^T - M_G^*)/M_G^*$. The absolute excess gain reports how much welfare PI traders collectively take away from TT traders with respect to the initial pie. The relative excess gain measures how much (on average) a PI trader is expected to improve his final welfare by trading within a given protocol.

A graphical representation of each set of data is obtained as follows. Given a protocol T, we plot the 999 data points produced in a simulation. We then fit a smoothing function generated by applying a Friedman smoother to all the data points associated with the same protocol; see Venables and Ripley (2002). Reading Figure 6.2 from left to right exemplifies this procedure for the case of a continuous double auction.
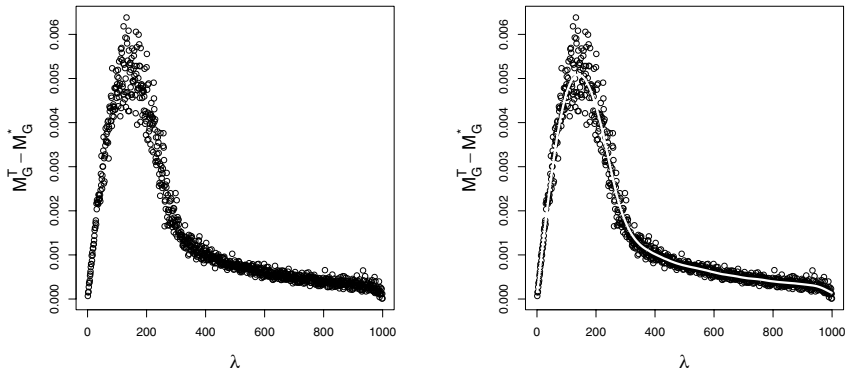

Fig. 6.2. Realizations (left) and a superimposed Friedman smoother (right).

6.4 Results

Figure 6.3 shows two representative pictures based on our exemplar. The figure on the left reports the (normalized) absolute excess gain

$M_G^T - M_G^*$ collectively achieved by the PI traders as a function of their cardinality λ for four protocols: batch auction, continuous double auction, nondiscretionary dealership, and the hybridization of these two latter protocols. Note that dividing λ by $n = 1000$ gives the proportion π of PI agents active in the market.
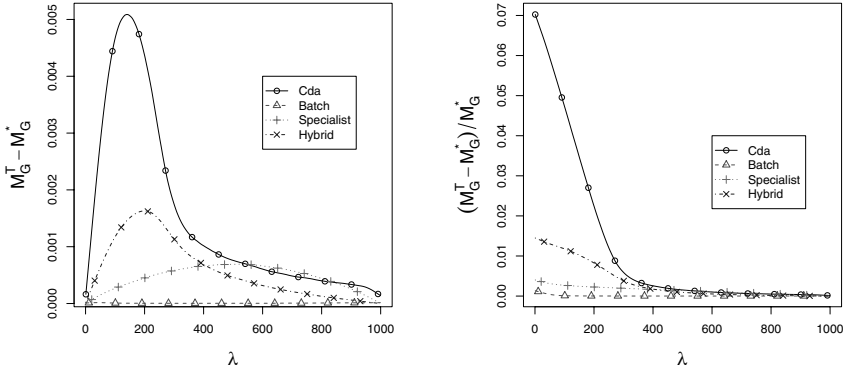
Fig. 6.3. Absolute (left) and relative (right) excess gain for PI traders.

The first comment is that the batch auction protects the TT traders much more effectively than any other protocol for both measures and for any number of PI traders. This is not surprising: the batch auction protocol requires simultaneous submission of trading orders and is therefore much more difficult for PI agents to exploit. By posting an order at p^*, each PI trader maximizes the probability of trading under the constraint of never losing on his fair share. Whenever the trading price issued in a session of the batch auction is different from p^*, he cuts away a piece of a TT trader's fair share. However, because the batch auction aggregates all the orders received in a trading session, it is very unlikely to issue a trading price different from p^*. We can thus shift our focus of interest to the three sequential protocols. For completeness, however, we report also the data relative to the batch auction.

The second comment is that in general the absolute excess gain for sequential protocols is a unimodal function of λ. Therefore, the collective ability of PI to exploit TT traders peaks at some intermediate value of λ. In this respect, there is a *natural ordering* of protocols from dealership to hybrid protocol to continuous double auction that appears

twice. First, the maximum excess gain for PI traders is increasing in the natural ordering of protocols. Simultaneously, the value of λ that maximizes the PI excess gain is decreasing. In other words, PI traders can achieve a greater excess gain in a continuous double auction, but their most effective proportion in such protocol is lower.

The result that the excess gain is increasing in the natural ordering is a direct consequence of the "protection" that the dealership provides. Because the dealer posts bid and ask prices that tend to straddle the correct p^*, the transaction price is never too different from this latter price; hence, no much fair share can be lost. The result that the λ's maximizing total excess gain are decreasing in the natural ordering can be heuristically explained by the combination of two effects. Intuitively, PI traders are most effective to exploit TT traders when their proportion is neither too low (there must be enough exploiters around) neither too high (there must be enough people to exploit). But we can put a bit more flesh on this explanation.

Consider the continuous double auction. The overall fair share for the TT group that PI traders can appropriate is roughly proportional to $(1 - \pi)$. On the other hand, taken as a group, the TT traders can lose a piece of their fair share only when one of them trades with a PI agent at a price different from p^*. The probability of a PI agent being matched for trade with a TT agent is roughly proportional to $\pi(1-\pi)$. Therefore, the excess gain appropriated by the PI group in the continuous double auction are approximately proportional to $\pi(1-\pi)^2$ and the maximum should be attained around $\hat{\pi} = 1/3$, corresponding to $\lambda = n\hat{\pi} = 333$ in our exemplar. The actual value is somewhat lower because some matchings between PI and TT agents do not lead to any trade.

Consider now the dealership. The overall fair share that PI traders can appropriate is still roughly proportional to $(1 - \pi)$. Moreover, because they can only trade with an impersonal dealer, the probability of a trade involving a PI agent is roughly proportional to the fraction π. Therefore, the excess gain for the PI group is now approximately proportional to $\pi(1 - \pi)$ and the maximum should be attained around $\hat{\pi} = 1/2$, corresponding to $\lambda = 500$ in our exemplar. As before, the exact value of the maximizer is affected by microstructural considerations that this heuristic argument does not capture. Finally, the corresponding values for the hybrid protocol are a convex combination of those of the parent protocols.

The third comment is that there are no important differences among sequential protocols when π (or λ) is sufficiently large, because there

are too few TT traders to be exploited. The overall fair share to be appropriated is roughly proportional to $(1 - \pi)$ and for large π there is simply too little to be taken away by the PI group. Moreover, markets with a high proportion of PI traders tend to exhibit a similar degree of allocative fairness because a PI agent never loses on his fair share. Therefore, we restrict the following comparisons to $\pi \leq 40\%$, corresponding to $\lambda \leq 400$ in our exemplar. For any proportion $\pi \leq 40\%$, the ranking over sequential mechanisms concerning their ability to prevent PI agents from eroding TT traders' fair shares is clear-cut and follows the natural ordering.

The right-hand side of Figure 6.3 reports the relative extra gain $(M_S^T - M_S^*)/M_S^*$ collectively achieved by the PI traders as a function of their number λ for the three sequential protocols. Unsurprisingly, this shows that increasing the number of exploiters makes their "looting" less effective for each protocol. Moreover, the ranking is again clearcut and follows the natural ordering. Finally, this effect is essentially unchanged in all the additional tests reported in the following section.

6.4.1 Tests of robustness

We have run some robustness tests by changing one parameter at a time in the exemplar. The first test looks at differences in the total endowment of stock, leading to a different equilibrium price p^*. The exemplar has a ratio $S/K = 2$ yielding $p^* = 1000 - 120(S/K) = 760$ and generates the left-hand side of Figure 6.3. We keep the same k_i for each trader i, but endow him with a different multiple of his original endowment s_i. This changes the ratio s_i/k_i and of course S/K as well. Figure 6.4 reports data when $S/K = 1$ (on the left) and $S/K = 3$ (on the right), corresponding respectively to a smaller and to a larger total endowment of stock. The equilibrium prices are now 880 and 640, respectively. We adjust the initial dealer's quotes accordingly, making sure that they always exactly straddle the equilibrium price.

Comparing the two figures from Figure 6.4 (as well as the right-hand side of Figure 6.3) shows that a larger stock endowment S increases the relative excess gain $(M_G^T - M_G^*)/M_G^*$ of PI for each protocol and each λ. We do not report the figures for the absolute excess gain to preserve space, but they exhibit a similar increasing effect. In fact, the following argument shows that, over our range of choices for S/K, an increasing relative excess gain implies an increasing absolute excess gain. The quantity M_G^* is roughly proportional to πM^*; in turn, M^* is increasing in (S/K) as far as $\tau \mu \geq (S/K)$ — as seen by differentiating (6.4) with respect to S. As this inequality holds for our choices of S/K,
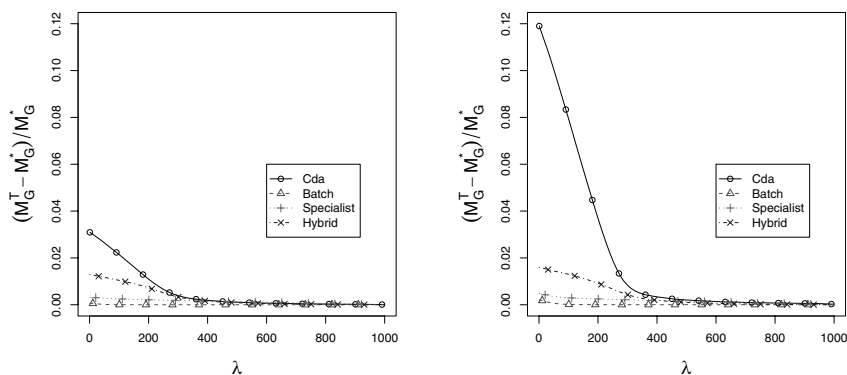
Fig. 6.4. Different equilibrium prices: $S/K = 1$ (left) and $S/K = 3$ (right).

the denominator of the relative excess gain is increasing and the claim follows. Put differently, this shows that a higher stock endowment S brings about a roughly proportional increase in the absolute excess gain.

The increase in the relative excess gain exhibited by Figure 6.4 is a stronger property that is explained by a second perhaps less obvious effect. *Ceteris paribus,* a larger S increases the number of trades that need to be carried out in order to reach the allocative efficiency. Each of these trades is a potential opportunity for PI agents to exploit, making them more likely to extract excess gain from the TT agents. This second effect accounts for the increase in the relative excess gain.

The second test considers the effect of changing the dealer's fixed spread in the nondiscretionary dealership, while keeping his initial quotes centered around the equilibrium price. The exemplar has a fixed spread of 10. Figure 6.5 reports the absolute excess gain when the fixed spread is 6 (bottom), 10 (middle), or 30 (top). The lower the spread, the more influential is the dealer's ability to constrain prices within a narrow band that individually rational trading naturally tends to keep around the equilibrium price p^*. Forcing the transaction price to lie in a band, of course, protects TT agents from more serious mispricings and hence reduces the ability of PI traders to exploit them. Accordingly, we see in Figure 6.5 that the absolute excess gain is increasing in the dealer's fixed spread for any number λ of PI traders.

A third test checks the effect of changing the dealer's fixed spread in the hybrid protocol dealership where an agent has access both to the dealer's quotes and to a book fed with limit orders from other traders.
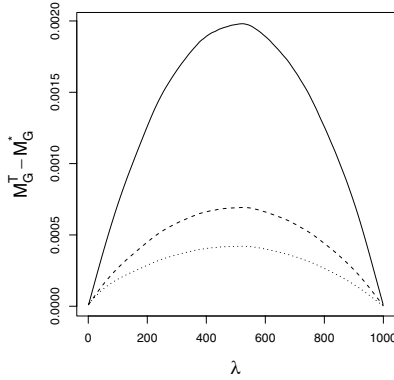
Fig. 6.5. Different fixed spreads in the dealership: 6, 10, 30 (bottom to top).

The exemplar has a fixed spread of 30. Figure 6.6 reports the absolute excess gain when this fixed spread takes five different values from 4 (bottom) to 300 (top). The absolute excess gain is increasing in the dealer's fixed spread for any number λ of PI traders. This effect and its explanation are analogous to the above. There is a second more interesting effect to note. In terms of the ability to control the absolute excess gain, the continuous double auction is the limit case of the hybrid protocol as the fixed spread goes to $+\infty$. When the dealer posts bid and ask that are too far apart, trading takes place only on the book. Accordingly, as we move from a low to a high spread, the excess gain curve morphs from the shape associated with a dealership to the shape associated with a continuous double auction; for instance, the peak increases and shifts leftward.

6.5 Conclusions

We have studied the performance of four market protocols with regard to their ability to equitably distribute the gains from trade among two groups of participants in an exchange economy. We assume Walrasian tatônemment as benchmark and define the fair share that should accrue to a trader as the certainty equivalent he would attain under this procedure.

When necessary, the first group of traders bids or asks their reservation value; this makes sure that trading never decreases their own certainty equivalent but exposes them to a possible loss on their fair
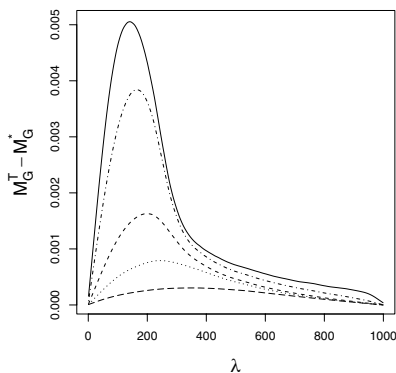
Fig. 6.6. Different dealer's spreads in the hybrid protocol: 4, 10, 30, 100, 300 (bottom to top).

share. The second group of traders knows (or can compute) the equilibrium price p^* and uses this information to make sure that trading cannot reduce either their certainty equivalent or their own fair share.

We test the allocative fairness of protocols by running (computerized) experiments where these two families of traders interact with each other. We find that there is a clear-cut ranking of protocols with respect to allocative fairness, defined as their ability to prevent PI agents from eroding TT traders' fair shares. Going from best to worst, this ranking is: batch auction, nondiscretionary dealership, the hybridization of a dealership and a continuous double auction, and finally the pure continuous double auction. The same ranking holds when we replace the absolute excess gain for PI traders with their relative excess gain.

References

E. Fehr, G. Kirchsteiger, and A. Riedl. Does fairness prevent market clearing? an experimental investigation. *The Quarterly Journal of Economics*, 108(2):437–459, May 1993.

D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual ra-

tionality. *Journal of Political Economy*, 101(1):119–137, February 1993.

L. Hurwicz. Economic design, adjustment and mechanism. *Economic Design*, 1:1–14, 1994.

M. LiCalzi and P. Pellizzari. The allocative effectiveness of market protocols under intelligent trading. In C. Bruun, editor, *Advances in Artificial Economics*, pages 17–29. Springer, 2006.

M. LiCalzi and P. Pellizzari. Simple market protocols for efficient risk sharing. *Journal of Economic Dynamics and Control*, 2007. Forthcoming.

W. N. Venables and B. D. Ripley. *Modern Applied Statistics with S.* Springer, fourth edition, 2002.

R. Wilson. The theory of syndicates. *Econometrica*, pages 119–132, 1968.