

Allocative Efficiency and Traders' Protection Under Zero Intelligence Behavior

Marco LiCalzi, Lucia Milone, and Paolo Pellizzari

Abstract This paper studies the continuous double auction from the point of view of market engineering: we tweak a resampling rule often used for this exchange protocol and search for an improved design. We assume zero intelligence trading as a lower bound for more robust behavioral rules and look at allocative efficiency, as well as three subordinate performance criteria: mean spread, cancellation rate, and traders' protection. This latter notion measures the ability of a protocol to help traders capture their share of the competitive equilibrium profits.

We consider two families of resampling rules and obtain the following results. Full resampling is not necessary to attain high allocative efficiency, but fine-tuning the resampling rate is important. The best allocative performances are similar across the two families. However, if the market designer adds any of the other three criteria as a subordinate goal, then a resampling rule based on a price band around the best quotes is superior.

1 Introduction

In a seminal paper, Gode and Sunder (1993a) define a *zero intelligence* (ZI) trader as an agent that “has no intelligence, does not seek or maximize profits, and does not observe, remember or learn.” (p. 121) Such zero intelligence assumption is not meant to provide a descriptive model of individual behavior: on the contrary, it is used to instantiate severe cognitive limitations that should impede the overall performance of the market.

A ZI agent is usually modeled as a robot player that submits random offers in an exchange market, under a minimal assumption of *individual rationality*: he never takes actions that can lead him to trade at prices below his cost or above his valuation. To the best of our knowledge, the first (unnamed) use of individually rational zero intelligence behavior in economic theory goes back to the B-process studied in Hurwicz et al. (1975); they prove that, if the market protocol allows unlimited

M. LiCalzi (✉)

Dept. of Applied Mathematics and Advanced School of Economics, University of Venice,
Dorsoduro 3825/E, 30123 Venice, Italy
e-mail: licalzi@unive.it

retrading, an economy without externalities must converge to a Pareto optimal allocation. Throughout this paper, we take the postulate of individual rationality for granted and speak simply of zero intelligence behavior.

By simulating the actions of (individually rational) ZI traders in a continuous double auction, Gode and Sunder (1993a) achieved levels of allocative efficiency similar to the outcomes generated by human subjects in laboratory experiments. This was used to argue that the main feature leading to a high allocative efficiency is the market protocol rather than the trading strategies used by the agents. More boldly put, the market can substitute for the cognitive limitations of the individuals. This conclusion has spawned a large literature venturing in different directions, including experimental economics and computer science; see Duffy (2006) for a thorough survey.

In general, it is widely acknowledged that the interpretation of Gode and Sunder's results is controversial. Gjerstad and Shachat (2007) emphasize the role of individual rationality as the key crucial assumption for allocative efficiency. A recurrent theme is the robustness of Gode and Sunder's conclusion: it is not difficult to produce environments where the allocative efficiency reached by ZI agents badly underperforms humans' results; see e.g. Brewer et al. (2002). On the other hand, the literature has shown that even minor improvements to the basic ZI trading rules suffice to achieve convergence to the competitive equilibrium; see Cliff and Bruten (1997) or Crockett et al. (2008).

Clearly, humans' cognitive abilities provide more leverage than zero intelligence. Therefore, we do not expect that the performance of a market protocol in an environment populated with ZI agents would be the same as with human traders. On the other hand, it is not unreasonable to postulate that the performance of a market protocol under a ZI behavioral assumption provides a plausible benchmark for its evaluation in view of use by human subjects. In his recent discussion of the "market-centered theory of computational economics", Mirowski (2007) attributes to the zero intelligence literature the computational insight that human cognitive abilities can be ignored under controlled circumstances to focus on the causal capacities of the market protocols. In a similar vein, Sunder (2004, p. 521) states that "[w]hen seen as human artifacts, a science of markets need not be built from the science of individual behavior." The implicit claim is that we may learn about the properties of markets regardless of the agents operating in them.

Our viewpoint is the following. Market protocols are complex artifacts; see Subrahmanian and Talukdar (2004). Their design requires a special attention to details and minutiae that partakes of the engineering attitude advocated in Roth (2002): we need to complement theory with experiments and computational simulations. In order to make fine-grained comparisons among different protocols, it is necessary to pin down agents' behavior to a simple standard. The ZI assumption provides a rough simplification under which it is possible to evaluate markets protocols *in silico* in order to select more promising designs.

The purpose of this paper is to exemplify this approach with regard to the continuous double auction. We replicate the results produced in Gode and Sunder (1993a) and show that they depend crucially on a subtle assumption about the market pro-

tol that has gone unnoticed in the literature. They write: "There are several variations of the double auction. We made three choices to simplify our implementation of the double auction. Each bid, ask, and transaction was valid for a single unit. *A transaction canceled any unaccepted bids and offers.* Finally, when a bid and a ask crossed, the transaction price was equal to the earlier of the two." (p. 122, emphasis added). As discussed below, the second emphasized assumption forces a frequent resampling of agents' quotes that is crucial (under zero intelligence) for allocative efficiency. We call this assumption *full resampling*: speaking figuratively, it mandates to toss away the book after each transaction. This seems both unrealistic and unpalatable for practical market design.

We are thus left to ask whether Gode and Sunder's implementation of the continuous double auction is a promising design. Taking the viewpoint of a market designer who is interested in allocative efficiency, we evaluate alternative market protocols that enforce different degrees of resampling. As it turns out, the assumption of full resampling is not necessary to achieve very high allocative efficiency under zero intelligence. There is a continuum of protocols, ordered by the strength of their resampling properties, that attain comparable levels of efficiency. This makes it possible to search for more effective protocols than Gode and Sunder's (1993a) without renouncing the objective of allocative efficiency.

To refine our selection, we introduce a subordinate criterion. While allocative efficiency is desirable from an aggregate point of view, a single trader in an exchange market is likely to be more interested in getting a fair deal. Let the *competitive share* of a trader be the profit he would make by transacting at the (competitive) equilibrium price. A market protocol that is more effective in helping traders realize their competitive share offers a superior *traders' protection*. Therefore, we study the traders' protection offered by comparably efficient market protocols to devise a practical and simple implementation of the continuous double auction.

We study two families of resampling rules and identify a design that delivers a significant improvement over Gode and Sunder's (1993a). However, barring an experimental validation with human subjects, we can only claim that the lower bounds on its performance with regard to both allocative efficiency and traders' protection are higher under zero intelligence.

The organization of the paper is the following. Section 2 describes the model used in our computational experiments and clarifies some technical details in the implementation of Gode and Sunder's (1993a) continuous double auction. The zero intelligence assumption is maintained throughout the paper. Section 3 proves that some (possibly not full) resampling is a necessary condition for allocative efficiency in the continuous double auction; see also LiCalzi and Pellizzari (2008). Section 4 shows that partial resampling may be sufficient for allocative efficiency. Based on this result, we study a family of resampling rules for the implementation of the continuous double auction protocol that modulates the probability of clearing the book after a transaction. Several rules within this family attain comparable levels of allocative efficiencies. Section 5 introduces an alternative way to effect resampling that is based on the use of a price band. Section 6 compares the alternatives and argues that the second method delivers a better protocol. Section 7 recapitulates our conclusions.

2 The Model

We use a setup very similar¹ to Gode and Sunder (1993a), who consider a simple exchange economy. Following Smith (1982), we identify three distinct components for our (simulated) exchange markets. The environment in Sect. 2.1 describes the general characteristics of our simulated economy, including agents' preferences and endowments. Section 2.2 specifies how agents make decisions and take actions under the zero intelligence assumption. This behavioral rule is kept fixed throughout this paper to let us concentrate on the effects of tweaking the market design. Finally, Sect. 2.3 gives a detailed description of the institutional details that form the protocol of a continuous double auction (and its variants) which regulate the exchange.

2.1 The Environment

There is an economy with a number n of traders, who can exchange single units of a generic good. (We set $n = 40, 200, 1000$ to look at size effects.) Each agent is initialized to be a seller or a buyer with equal probability. Each seller i is endowed with one unit of the good for which he has a private cost c_i that is independently drawn from the uniform distribution on $[0, 1]$. Each buyer j holds no units and has a private valuation v_j for one unit of the good that is independently drawn from the uniform distribution on $[0, 1]$. Without loss of generality, prices are assumed to lie in $[0, 1]$.

2.2 Zero Intelligence Behavior

Zero intelligence reduces behavior to a very simple rule: when requested a quote for an order, a trader draws a price from a random distribution (usually taken to be uniform). We assume that traders' behavior abides by *individual rationality*: each seller i is willing to sell his unit at a price $p \geq c_i$ and each buyer j is willing to buy one unit at a price $p \leq v_j$. Therefore, throughout this paper, the zero intelligence assumption pins down behavior as follows: when requested a quote for an order, a seller i provides an ask price that is an independent draw from the uniform distribution on $[c_i, 1]$; similarly, a buyer j makes a bid that is an independent draw from the uniform distribution on $[0, v_j]$. This behavioral rule is called ZI-C in Gode and Sunder (1993a).

Note that the only action requested by an agent is to issue a quote: it is left to the market to process traders' quotes and execute transactions on their behalf.

¹There are negligible differences. We consider n agents who can trade at most one unit, while they have 12 traders who can exchange several units but must trade them one by one. Our setup is simpler to describe because it associates with each trader a single unit and a one-dimensional type (his cost/valuation).

This is consistent with an approach of market engineering: we are not interested in the performance of more sophisticated behavioral rules, but rather in the design of protocols that take decent care even of simple-minded agents.

In particular, this paper studies protocols that implement variants of the continuous double auction, where agents sequentially place quotes on the selling and buying books. Orders are immediately executed at the outstanding price if they are marketable; otherwise, they are recorded on the books with the usual price-time priority and remain valid unless a cancellation occurs. When a transaction takes place, the orders are removed from the market and the traders leave the market and become inactive.

The zero intelligence assumption places a second restriction on agents' behavior. In a sequential protocol like the continuous double auction, an agent can choose both his action and the time at which to take it; see Gul and Lundholm (1995). Zero intelligence robs agents of the opportunity to make decisions about the timing at which to issue a quote. Agents are exogenously arranged in a queue and reach the market one at a time, until the queue is exhausted or some exogenous event triggers the formation of a new queue.

The standard implementation is the following. At the beginning of a simulation, all agents are active and placed in the queue. If an agent reaches the market and trades his unit, he becomes inactive for the rest of the simulation. Otherwise, he is in one of two states: either he has an order on the book (because he is active and the queue has already reached him), or he is still queueing for a chance to act. An important detail in the design of an experiment is the set of events that triggers the formation of a new queue, reshuffling the state of active agents. For instance, the full resampling assumption in Gode and Sunder (1993a) makes each transaction a trigger event that sends all traders with an order on the book back to the end of the queue.

2.3 The Protocol

The implementation of the continuous double auction in Gode and Sunder (1993a) is based on several rules. Some of them are not stated explicitly in the paper, but may be gathered by a joint reading of other related papers; see in particular Gode and Sunder (1993b, 2004). For completeness and ease of reference, we collect here all the ingredients we found necessary to replicate their results.

The first three rules correspond to the assumptions cited above. We begin with the first and the third. The *single unit* rule states that all quotes and prices refer to one unit of the good. A standard rule of *precedence* decides the transaction price: when two quotes cross, the price is set by the earlier quote. We maintain both the single unit and the precedence rules, because they entail no loss of generality.

The second of the three assumptions put forth in Gode and Sunder (1993a) as "simplifications" states that the book is cleared after each transaction. By itself, this rule is surprising because tossing away the book at any opportunity seems to run

contrary to the obvious purpose of storing past orders and make them available to future traders. In Gode and Sunder's design, moreover, this rule triggers a refreshing of the queue: after each transaction, all recorded orders are deleted and their owners are given a new chance to act. When a ZI agent goes back to the queue and comes up again, he randomly issues a new quote. Hence, the real consequence of tossing away the book is to free up the agents' past quotes and force them to issue novel ones. That is, after each trade, all active agents who have placed an order since the former transaction are resampled. This is the reason for calling their assumption *full resampling*. Section 3 shows that full resampling is crucial for Gode and Sunder's results and hence cannot be dismissed as a mere "simplification". In fact, one of the motivations for this paper is to study the import of this neglected assumption.

There are other rules that need to be made explicit. *No retrading* states that buyers and sellers can never exchange roles: a buyer (seller) who acquires (transfers) a unit is not allowed to sell (buy) it later to other traders. The intuition that, given sufficient retrading, a market populated with ZI agents should reach full allocative efficiency is proven in Hurwicz et al. (1975). Therefore, no retrading is necessary to avoid trivialities. Gode and Sunder (2004) provide further comments on the role and plausibility of this assumption.

The *uniform sequencing* of agents within a simulation arranges them in a queue according to an exogenously given order, which is independently drawn from the uniform distribution over all permutations of agents. As explained in Gode and Sunder (2004), in their simulations the queue of traders is sampled *without replacement*. That is, when the execution of a transaction triggers a refreshing of the queue, the agents who have a quote stored on the book re-enter it behind the traders still waiting in the original queue. The no replacement assumption is a sensible simplification that allows for faster and more efficient coding. However, since this rule violates anonymity, its practical implementation requires either additional information processing (when control is centralized) or some traders' coordination (under decentralization). Therefore, in the interest of simplicity and realism, we maintain the uniform sequencing rule but we switch to sampling with replacement: when an event triggers the formation of a queue, we simply apply uniform sequencing over all active agents.

Finally, the *halting* rule mandates when a trading session is over. In Gode and Sunder (1993a) a trading session is a period of fixed duration that lasts 30 seconds for each computational simulation. Traders are put in a queue and asked to provide a quote. If all the queued agents have issued a quote and no transaction has occurred, the books are cleared and a new queue is started until time is over. Given that robot players are remarkably fast, this implies that an agent is likely to be asked to issue a quote several times. (We have been unable to determine how often queues are restarted in Gode and Sunder (1993a) only because the time limit has not been reached.) Unfortunately, given that hardware (and software) vary in processing power (and efficiency), a halting rule based on a fixed duration is not sufficient to ensure comparable results when simulations are run on different machines. Therefore, we choose a different halting rule that allows for full comparability: a trading session is over when the queue of traders waiting to place an order is exhausted.

An additional advantage of this assumption is that it biases our simulations in the proper direction: *ceteris paribus*, we resample traders less often because our halting rule is more stringent. This makes allocative efficiency harder to attain.

For the reader's convenience, we recap here the rules of the continuous double auction protocol used in all the simulations discussed in this paper. We use single unit trading, set the transaction price by precedence, exclude retrading, and apply uniform sequencing. Differently from Gode and Sunder (1993a), we put traders back in the queue with replacement and use a more restrictive halting rule.

3 The Resampling Assumption

We test the import of the resampling assumption for the allocative efficiency of the continuous double auction (as implemented by our protocol). As usual, we define allocative efficiency as the ratio between the realized gains from the trade and the maximum feasible gains from trade, which can be formally defined as done in Zhan et al. (2002, p. 678). This measure is adimensional, facilitating comparisons throughout the paper.

3.1 Resampling Is Necessary for Allocative Efficiency

We contrast full resampling against no resampling under zero intelligence trading. *Full resampling* mandates that after each transaction the book is cleared and active traders with an order on the book are sent back to the waiting queue. *No resampling* postulates that submitted orders stay on the book until the end of the trading session (unless they are used up for a transaction); e.g. see Maslov (2000).

The difference between full and no resampling is stark. A ZI agent acts only when its turn in the waiting queue comes up. Under no resampling, each agent is given only one chance to act by sending a random quote to the book. Under full resampling, on the other hand, until an agent completes a trade and becomes inactive, any refresh of the waiting queue following a transaction may give him a new chance to act and generate another random quote. Therefore, the number of opportunities for actions is much greater under full resampling, and this should increase allocative efficiency.

The datapoints on the left-hand side of Fig. 1 represent the allocative efficiencies under *full resampling* for 500 different runs with $n = 200$ agents. The data match Gode and Sunder's (1993a) results, confirming that the impact of our (more stringent) halting rule on allocative efficiency is negligible. The right-hand side provides analogous information for the case of *no resampling*. The y-axes use the same scale, so that a direct comparison by visual inspection is immediate: the higher the level, the higher the allocative efficiency.

The difference in performance under full or no resampling is remarkably substantial. The mean (median) allocative efficiency is 0.910 (0.916) with full resampling and 0.497 (0.498) with no resampling. (All statistics reported in this paper




Fig. 1 Allocative efficiency under full (*left*) or no resampling (*right*)

are rounded to the closest third decimal digit.) The min–max range (standard deviation) for the allocative efficiency is $[0.759, 0.978]$ (0.036) with full resampling and $[0.246, 0.684]$ (0.070) with no resampling. Within our sample, the worst allocative efficiency with $n = 200$ agents under full resampling (0.759) is much higher than the best allocative efficiency under no resampling (0.684). Visual inspection strongly suggests that the distribution of the allocative efficiency under full resampling stochastically dominates the distribution under no resampling.² More modestly, we claim that the expected value of the allocative efficiency under full resampling is higher. In fact, the Wilcoxon signed-rank test rejects the hypothesis that the means are equal at a level of significance of 10^{-3} . (Throughout the rest of the paper, unless otherwise noted, we use the Wilcoxon signed-rank test to compare means and we require a p -value lower than 10^{-3} to claim statistical significance.)

Similar effects occur for different values of n , but a larger number of agents tends to improve allocative efficiency. Thus, when comparing data for a different number of agents, we should take into account a fixed size effect. We believe that $n = 200$ is a representative case, but for comparability Table 1 lists the main statistics for $n = 200/5 = 40$ and $n = 200 \times 5 = 1000$.

Based on the relative size of the agents' pool, we say that the market is *thin* ($n = 40$), *thick* ($n = 200$), or *crowded* ($n = 1000$). Each column summarizes 500 distinct simulation rounds.

It is apparent that no resampling may be calamitous in a thin market, because an agent who happens to issue a “wrong” quote is given no later chance to remedy. Analogously, a few “lucky” trades may shoot allocative efficiency up. Hence, the dispersion of the allocative efficiency is much higher in a thin market. Such effects are washed out in a crowded market. Overall, an increase in n has a positive effect on allocative efficiency under either resampling assumption. But the effect is sharper under full resampling, because this rule gives traders more chances to trade.

²LiCalzi and Pellizzari (2008) document a similar effect over four different trading protocols.

Table 1 Summary statistics for the allocative efficiency

	Full resampling			No resampling		
	$n = 40$	$n = 200$	$n = 1000$	$n = 40$	$n = 200$	$n = 1000$
Mean	0.735	0.910	0.949	0.441	0.497	0.517
Median	0.765	0.916	0.951	0.456	0.498	0.518
Minimum	0.053	0.759	0.911	0.000	0.246	0.405
Maximum	1.000	0.978	0.971	0.933	0.684	0.609
Std. dev.	0.169	0.036	0.009	0.157	0.070	0.032

Our experiment shows that, *ceteris paribus*, full resampling yields a much higher allocative efficiency than no resampling. Speaking figuratively, no resampling switches off the ability of a protocol to help ZI agents capture most of the available gains from trade. We conclude that (at least some) resampling is a necessary condition for allocative efficiency. This reduces the scope of Gode and Sunder's (1993a) results about the ability of a market to substitute for agents' lack of rationality: an effective protocol for ZI agents must include rules that ensure an adequate amount of resampling.

On the other hand, our results do not invalidate their claim that it is possible to design markets that may overcome agents' cognitive limitations. To the contrary, they suggest that the use of a (partial) resampling rule may be a particularly clever design choice for fostering allocative efficiency in exchange markets. Section 4 sets out to examine a continuum of alternative rules that enforce different degrees of resampling in this respect. We find that less than full resampling is sufficient to reach high levels of efficiency.

3.2 Efficiency and Full Resampling

Before moving to issues of market engineering, there are two hanging questions to address. First: why does full resampling lead to higher allocative efficiency than no resampling? Second: where does the efficiency loss go?

We begin with the second question, whose answer leads naturally to the first one. Gode and Sunder (1997, p. 605) point out that in general there are "three causes of inefficiency: (1) traders participate in unprofitable trades; (2) traders fail to negotiate profitable trades' and (3) extramarginal traders displace intramarginal traders." Since individual rationality rules out the first source of inefficiency, we need be concerned only with the other two. They can be measured; e.g., see Zhan and Friedman (2007) who also provide formal definitions.

Let p^* be the market-clearing price. (There may be an interval of market-clearing prices. We assume that p^* is the midpoint.) Individually rational traders who would transact at p^* are called *intramarginal*; all other traders are *extramarginal*. If at the end of a trading session an intramarginal trader i has failed to trade, this creates a

Table 2 A breakdown of the efficiency loss

	Full resampling			No resampling		
	$n = 40$	$n = 200$	$n = 1000$	$n = 40$	$n = 200$	$n = 1000$
<i>AE</i>	0.735	0.910	0.950	0.441	0.497	0.517
<i>MT</i>	0.241	0.055	0.012	0.548	0.495	0.477
<i>EM</i>	0.025	0.035	0.037	0.011	0.008	0.006

loss of total surplus equal to $v_i - p^*$ if he is a buyer and $p^* - c_i$ if he is a seller. The sum of these losses corresponds to (2) above: we call it *MT*, as a mnemonic for the inefficiency caused by *missed trades*. The third case comes about when a transaction involves an extramarginal trader, causing a loss equal to his profit at p^* . The sum of such losses corresponds to (3) above: we call it *EM*, as a mnemonic for the inefficiency due to *extramarginal* trades. As discussed in Zhan and Friedman (2007), the allocative efficiency decomposes as $AE = 1 - MT - EM$; or, equivalently, $MT + EM = 1 - AE$ measures the allocative inefficiency. Table 2 provides a breakdown of the efficiency loss for thin, thick and crowded markets by listing mean values over 500 distinct simulation rounds. Values may not add up to 1 because of rounding effects.

There are two observations to be made. The first one is that the efficiency loss (*MT*) attributable to missed trades is decreasing in the thickness of the market, because thicker markets facilitate the search for a matching quote. Moreover, trading under no resampling terminates too soon: most of the efficiency loss comes from missed trades. (The difference between the mean values for *MT* under full or no resampling is statistically significant.) The reason for a high allocative efficiency under full resampling is elementary: this rule is of course more effective in prolonging the trading session, and hence gives traders enough chances to find their right match. This suggests that an effective market protocol should offer agents an adequate number of matching opportunities to keep the *MT* component of the efficiency loss under control.

The second observation points out a shortcoming of full resampling. The average value of *EM* is higher under such rule. (The difference between the means is once again statistically significant.) This is not difficult to explain: by the precedence rule, the best outstanding bid and ask in the book bound the price of the next transaction. The narrower the spread, the more difficult is to steal a deal for an extramarginal trader. Storing earlier quotes in the book provides (intramarginal) traders with some price protection and makes them less exploitable by extramarginal agents. As the full resampling rule tosses away the book after each transaction, it renounces such protection all too frequently (compared to no resampling). This is apparent by a straightforward comparison: the average spread (sampled before a trader places an order) is 0.152 with no resampling and 0.327 with full resampling when $n = 200$. (Corresponding values are 0.267 and 0.398 for $n = 40$; 0.113 and 0.268 for $n = 1000$.) The differences between the mean values are statistically significant. Since the zero intelligence assumption prevents traders from adjusting their quotes




Fig. 2 Allocative efficiency under π -resampling

based on the state of the book, full resampling is a rule more favorable to extra-marginal traders than no resampling. This suggests that an effective market protocol should incorporate some form of price protection to keep the *EM* component of the efficiency loss under control.

4 Randomized Resampling

Section 3.1 established that the resampling rule is crucial to reach allocative efficiency under zero intelligence. To evaluate its impact, this section begins by looking at a continuum of resampling rules that generalize the simple dichotomy between no and full resampling. We emphasize that these rules are chosen to compare and understand how resampling affects the trading protocol. Like engineers, we are searching for improvements and tweaks over a basic design.

4.1 Full Resampling Is Not Necessary for Allocative Efficiency

A simple way to conceptualize the distinction between no and full resampling is to note that these two rules react differently to the same event; namely, the occurrence of a transaction. When two orders cross, full resampling clears the book with probability one whereas no resampling does so with probability zero. This naturally suggests to consider a family of randomized resampling rules that clear the book with probability π in $[0, 1]$ whenever there is a transaction. This set embeds full resampling for $\pi = 1$ and no resampling for $\pi = 0$.

The right-hand side of Fig. 2 shows the allocative efficiency under π -resampling with $n = 200$ agents. The graph is obtained as follows. We choose the 21 equispaced points $\{0, 0.05, 0.10, \dots, 0.90, 0.95, 1\}$ in the $[0, 1]$ interval. For each of these π -values, we run 500 distinct simulations. The allocative efficiencies obtained over

these $21 \times 500 = 10500$ simulations are plotted as datapoints on the left-hand side of Fig. 2. We summarize these data by the mean allocative efficiency for each π . (The difference between a mean and the corresponding median is never greater than 0.022.) The 21 sample averages are joined using segments to obtain the thicker central piecewise linear representation.³ The two external thin graphs are similarly obtained by joining respectively the minimum and maximum values obtained for the allocative efficiency at a given value of the resampling probability π . We emphasize that the resulting band is not a confidence interval but the actual range of efficiencies obtained under our simulations: its main purpose is to provide a simple visual diagnostic for the dispersion of the data around their central tendency. We adopt the usual $[0, 1]$ -scale for the y-axis.

The graph on the right of Fig. 2 is easily interpreted. As expected, allocative efficiency is on average increasing in the probability π that a transaction triggers a clearing of the book. Under zero intelligence, the frequency with which resampling takes place has a direct effect on the ability of the protocol to reap high levels of efficiency. On the other hand, the graph shows also that full resampling ($\pi = 1$) is not necessary: the (average) allocative efficiency in our simulations is more than 90% for $\pi \geq 0.7$ with a (statistically insignificant) peak of 91.08% at $\pi = 0.95$; the standard deviations are never greater than 0.132. There is an upper bound on the allocative efficiency that can be attained but a sufficiently large π is enough to approach it.

Similar results hold for thin and crowded markets: when $n = 40$, $AE \geq 67\%$ for $\pi \geq 0.7$ with a peak of 73.48% at $\pi = 1$ and standard deviations never greater than 0.209; when $n = 1000$, $AE \geq 93\%$ for $\pi \geq 0.2$ with a (statistically insignificant) peak of 95.12% at $\pi = 0.85$ and standard deviations never greater than 0.080. The thickness of the market affects the upper bound on the allocative efficiency but, in general, there is a whole range of resampling probabilities that achieve comparably high levels of allocative efficiency under zero intelligence.

Our conclusion is that full resampling is not necessary for allocative efficiency. Full resampling sets $\pi = 1$ and tosses the book away after each transaction: this yields a high allocative efficiency under zero intelligence, but it is also an extreme assumption that is likely to be unpalatable for human traders in real markets. As it turns out, we can temper the strength of full resampling at the mere cost of a tiny reduction (if any) in allocative efficiency.

This leads naturally to frame the choice of a resampling rule as a tradeoff between its allocative benefits and its implementation costs. On the part of the market designer, there are obvious costs to continuously monitor and update the state of the book. Similarly, traders who are forced to check whether their past orders have been voided are likely to resist frequent cancellations. Intuitively, when the costs of full resampling are not trivial, we expect partial resampling ($0 < \pi < 1$) to be preferable. The rest of this section fleshes up this argument. Section 5 takes up a related

³We consistently apply this approach to construct the graphs for this paper: a broken line joins 21 points, each of which represents a statistic over 500 distinct simulations for a fixed value of a parameter such as π .

question and examines a different family of resampling rules to find out whether they perform better than π -resampling.

4.2 Where Is the Best π ?

Let us take stock of the starting point we have reached so far. First, given the thickness of the market, there is an upper bound on the (mean) allocative efficiency that can be attained using π -resampling. Second, the set of π -values for which the protocol reaches comparably high levels of efficiency is an interval. Thus, we need to look at additional performance criteria in order to pinpoint a smaller interval for the choice for π .

We do not claim that it is possible to find the *best* π and reduce such interval to a singleton, because the zero intelligence assumption provides at best a lower bound for the evaluation of a protocol. More modestly, we can define plausible performance criteria and measure them for different values of π under zero intelligence trading. Clearly, this procedure cannot provide a final verdict for the performance of the protocol with human subjects. Hence, the aim of this section is to carry out an engineering exercise and derive a robust choice: what is the range of π for which performance under zero intelligence is better, and why?

We consider two simple criteria. (Others are of course possible, and we take up a third major one in Sect. 4.3.) The first criterion deals with the basic requirement that an effective market protocol should offer some guidance to traders' choice in the form of a price signal. The closer the outstanding bid and ask straddle the (competitive) equilibrium price, the stronger the information that they provide. It is obvious that zero intelligence makes no use of this information: therefore, the object of our investigation is the ability of the protocol to provide an effective price signal *independently* of traders' behavior.

We measure it by the (mean) spread on the market: the closer the spread, the stronger the signal. The average is taken by sampling data when a trader arrives and places an order (as opposed to just before a transaction occurs), because we are interested in the state of the book found by a generic agent reaching the market. As it turns out, our environment is sufficiently regular that the best bid and the best ask are (on average) symmetric around the equilibrium price p^* . Hence, the outstanding spread is a sufficient statistic for such purpose. The left-hand side of Fig. 3 shows the (mean) outstanding bid and ask under π -resampling. The y-axis is truncated to [0.3, 0.7] to enhance readability.

Unsurprisingly, the spread is on average increasing in π . When resampling is more frequent, the book is cleared more often and hence is more likely to have both fewer quotes and a larger spread. For $n = 200$, the average (median) spread in our simulations increases monotonically from 0.152 (0.124) at $\pi = 0$ to a peak of 0.327 (0.232) at $\pi = 1$; the standard deviations are never greater than 0.033. Qualitatively similar results hold for $n = 40$ and $n = 1000$, and spreads are smaller in thicker markets. This leads to the following general piece of advice. Suppose




Fig. 3 Mean spreads and cancellation rates under π -resampling

that, conditional on achieving comparable levels of allocative efficiency, a market designer prefers narrower spreads. Then he should aim towards choosing a level of π that is bounded away from zero (to achieve efficiency) as well as from one (to obtain smaller spreads). The thicker the market, the weaker the need to stay away from one.

A second simple criterion has to do with the number of cancellations imposed on traders. (Recall that traders cannot cancel their orders.) The benefit of a cancellation is to offer a new chance for action to the trader. On the other hand, in general there are costs associated with the inconvenience of monitoring the state of an order or placing a new one. Therefore, when the allocative efficiency of two protocols are similar, it is reasonable to expect that the one leading to fewer cancellations should be preferred. We measure the cancellation rate as the average of the ratio between the number of orders canceled over the number of transactions completed over each of our 500 simulated trading sessions. Clearly, allocative efficiency is strongly correlated with volume; hence, the higher the ratio, the higher the cost of redundant cancellations. The right-hand side of Fig. 3 depicts the (mean) cancellation ratio with $n = 200$ agents. As usual, we report the mean values as a thick black line surrounded by thinner red lines that correspond to the minima and maxima.

Similarly to the spread, the cancellation rate is on average increasing in π because the amount of resampling directly correlates with the number of canceled orders. For $n = 200$, the mean (and standard deviation) of the cancellation rate go up⁴ from 2.643 (2.498) at $\pi = 0.05$ to a peak of 18.22 (3.002) at $\pi = 1$; the standard deviations are never greater than 3.124. Similar results hold for $n = 40$ and $n = 1000$, and mean cancellation rates are higher in thicker markets. The conclusion we draw is similar to the earlier one. Suppose that, conditional on achieving comparable levels of allocative efficiency, a market designer prefers a lower cancellation rate. Then an optimal π should be bounded away from zero (for efficiency)

⁴We start from $\pi = 0.05$ because the cancellation rate at $\pi = 0$ is zero by assumption.

as well as from one (for a lower rate). The thicker the market, the stronger the need to stay away from one.

4.3 Traders' Protection

The last performance criterion that we consider in this paper is directly inspired by Stigler (1964), who pioneered the use of simulations to address issues of market engineering. He put down a clear statement: "The paramount goal of the regulations in the security markets is to protect the innocent (but avaricious) investor" (p. 120). While his paper is concerned with security markets, the conditions for achieving this goal should also be investigated for exchange markets. Curiously, the literature on zero intelligence has so far neglected this issue to the point that there is not even an agreed convention on the exact meaning of protection.

This section provides a measurable criterion for traders' protection in an exchange market, and then applies it to the evaluation of the π -resampling rule. Ideally, in a competitive equilibrium, all⁵ the intramarginal traders exchange the good at the same equilibrium price p^* : nobody pays (or is paid) differently from the others. On the other hand, a continuous double auction offers neither of these guarantees: first, an intramarginal trader may fail to close a deal; second, the price at which a trade occurs may be different from the price agreed for another trade. Both of these events deny the competitive outcome to the intramarginal trader. When a market protocol holds such events under control, it manages to offer *traders' protection*.

Clearly, allocative efficiency does not measure traders' protection: since it focuses on the gains from trade that are realized, it fails to register at what terms these gains materialize. We need a more sophisticated measure that takes into account the price at which a transaction is carried out, and hence touches on the distribution of gains. To this purpose, we define the *competitive share* of a trader as the (positive part of the) profit he would make by transacting at the competitive equilibrium price. Given an equilibrium price p^* , the competitive share of a buyer with valuation v is $(v - p^*)^+$ and that of a seller with cost c is $(p^* - c)^+$. Clearly, the competitive share of any extramarginal trader is zero.

The *realized competitive share* is the portion of his competitive share realized by an agent. (Extramarginal traders are entitled to no competitive share.) If an agent fails to trade, this portion is zero. If a trade occurs at price p , the realized competitive share is $v - \max\{p, p^*\}$ for an (intramarginal) buyer and $\min\{p, p^*\} - c$ for an (intramarginal) seller.

The realized competitive share is concerned only with measuring whether a trader gets its due, and ignores any additional gains that he may be able to reap. The profit realized by an intramarginal trader may be greater than his realized competitive share if he manages to secure terms of trade more favorable than p^* ; similarly,

⁵When the number of intramarginal traders is odd, one of them will not trade for lack of a partner.




Fig. 4 Traders' protection (*left*), superimposed to allocative efficiency (*right*)

any extramarginal agent who completes a trade makes positive profits by individual rationality, but his realized competitive share remains zero.

Note that the sum of all the competitive shares equals the maximum feasible gains from trade. In analogy with allocative efficiency (AE), we define the traders' protection (for short, *TP*) offered by a market protocol as the ratio of the realized competitive shares and the sum of all the competitive shares. This measure is adimensional and takes values in $[0, 1]$.

The left-hand side of Fig. 4 shows the traders' protection under π -resampling with $n = 200$ agents. As usual, we report the mean values surrounded by minima and maxima. The right-hand side superimposes AE and TP to allow for a direct comparison: the black line corresponding to TP is the same visible on the left, while the red line depicting AE corresponds to the inner black line from the right-hand side of Fig. 2.

In general, traders' protection is not increasing in π . For $n = 200$, the mean protection starts at 0.431 in $\pi = 0$, peaks at 0.718 in $\pi = 0.7$ and then declines to 0.709 in $\pi = 1$ (with two local maxima of 0.711 at $\pi = 0.4$ and 0.715 at $\pi = 0.9$); the standard deviations are never greater than 0.111. Qualitatively similar results hold for crowded and thin markets. When $n = 1000$, TP is 0.461 in $\pi = 0$, peaks at 0.791 in $\pi = 0.2$ and then declines to 0.744 in $\pi = 1$ with no other local maxima and standard deviations never greater than 0.070. For $n = 40$, TP is 0.355 in $\pi = 0$ and peaks at 0.563 in $\pi = 1$, with four more local maxima in between and standard deviations never greater than 0.172.

Here, the thickness of the market has a very strong effect on the range of the best value for π : the more crowded the market, the smaller the resampling rate that provides the best protection. The overall conclusion is similar to the above, with a strong word of caution as regards the thickness of the market. Suppose that, conditional on achieving comparable levels of allocative efficiency, a market designer prefers to offer a higher traders' protection. Then an optimal π should be bounded away from zero (for efficiency) as well as from one (for protection). When the market gets thicker, however, the need to stay away from one is remarkably higher.

To sum it up, we have considered three criteria based respectively on spread, cancellation rate, and traders' protection. To a different extent, they all support a choice of π from the interior of the interval $[0, 1]$: at the cost of a nominal reduction in allocative efficiency, it is possible to have lower spreads, fewer cancellations, and higher traders' protection. It is clear that both the relative importance of these criteria to the market designer as well as the thickness of the market matter for the exact choice of π . However, generally speaking, all of our performance criteria strongly suggest that full resampling is unlikely to be a defensible choice.

5 Resampling Outside of a Price Band

Section 4 has studied randomized resampling, but it is obvious that there exist many other rules. It may be impossible to pick a best one, but we can compare the performance of different resampling techniques. This section considers a different rule that shares a few basic properties with π -resampling. First, it depends on a single parameter γ in $[0, 1]$. Second, it implies an average resampling rate that is increasing in the parameter. Third, it embeds the two extreme cases of full and no resampling for $\gamma = 1$ and $\gamma = 0$. Fourth, it requires minimal information and thus imposes very little burden on the market protocol or the cognitive abilities of the traders.

The γ -resampling rule is the following. After a trade carries out at price p , the protocol cancels all outstanding orders that fall outside the *price band* $[\gamma p, \gamma p + (1 - \gamma)]$; moreover, in the special case $\gamma = 1$, we require the protocol to erase even the outstanding orders at price p so that it clears the book entirely. (This specification is necessary to embed full resampling, because the book might contain orders with price p but lower time priority.) It is useful to keep in mind that π is the probability with which the book is cleared after a transaction, while $(1 - \gamma)$ is the width of the price band within which orders are *not* deleted after a transaction.

Like π -resampling, the γ -resampling rule is triggered whenever a transaction occurs. Differently from it, its application implies that traders whose orders are deleted may infer a one-sided bound for the last transaction price. For instance, given γ , when a buyer sees that his past order at price p has been canceled, he can deduce that the last transaction price must have been strictly greater than p/γ . We do not view this a significant limitation, since it seems highly plausible that all agents would be given public access to such information. On the other hand, since it makes no use of the best outstanding bid and ask, the γ -resampling rule does not require to divulge this kind of information. This may be an additional advantage in view of the results in Arifovic and Ledyard (2007), who consider a sequence of call markets and show that the closed book design⁶ brings about a higher allocative efficiency than an open book in environments populated with human subjects or (non ZI) simulated agents. If similar results should suggest adopting a closed book design for

⁶In a closed book, traders learn only the clearing price after each call; in an open book, they are also told the quotes processed in that call.




Fig. 5 Allocative efficiency under π - and γ -resampling

the continuous double auction, both π - and γ -resampling are compatible. For the current study, it suffices to say that the ZI assumption precludes a direct comparison between closed and open book design, because it prevents agents from making use of any disclosed information.

The γ -resampling rule may be easily adapted in other dimensions. For instance, our definition embeds a symmetry assumption that may be removed. We choose the endpoints of the price band at the same distance from the extremes of the price range: the left endpoint is a convex combination between the last transaction price p and the minimum possible price, while the right endpoint is a convex combination between p and the maximum possible price. Clearly, this choice requires the implicit assumption that we know that p lies in the interval $[0, 1]$. More generally, when no bounds for the price are known, it suffices to set the price band to be the interval $[\gamma p, (1/\gamma)p]$ for γ in $[0, 1]$ or other analogous formulations.

Figure 5 shows the allocative efficiency under π -resampling (on the left) and γ -resampling (on the right) with $n = 200$ agents.

The graph on the left is the same as in Fig. 2. The graph on the right is the analog for γ -resampling: a thick inner black line joins the mean values, and two thin outer red lines join the corresponding minima and maxima. The directionality of the graphs is aligned because they depict two resampling rules that coincide for $\pi = \gamma = 0$ and $\pi = \gamma = 1$.

Both resampling rules are on average increasing in the corresponding parameter. However, the qualitative behavior is different. Under π -resampling, allocative efficiency picks up fast and rapidly settles on a plateau: for $n = 200$, the sample average is greater than 0.90 for $\pi \geq 0.7$. As already discussed, even moderate levels of π suffice to attain an adequate level of efficiency. On the other hand, efficiency under γ -resampling grows up more slowly and, quite interestingly, peaks at $\gamma < 1$: for $n = 200$, the sample average is greater than 0.90 for $\pi \geq 0.75$ and peaks at 0.923 for $\gamma = 0.85$; the standard deviations are never greater than 0.073.

Similar results hold for the case of thin and crowded markets. For $n = 40$, the sample average is greater than 0.67 for $\gamma \geq 0.65$ and peaks at 0.743 for $\gamma = 0.9$,

Table 3 Maximum allocative efficiency under π - and γ -resampling

	π -resampling			γ -resampling		
	$n = 40$	$n = 200$	$n = 1000$	$n = 40$	$n = 200$	$n = 1000$
Maximum AE	0.735	0.911	0.951	0.743	0.923	0.960
Maximizer (π, γ)	1.000	0.950	0.850	0.900	0.850	0.800

with standard deviations never greater than 0.197; for $n = 1000$, the sample average is greater than 0.94 for $\gamma \geq 0.7$ and peaks at 0.96 for $\gamma = 0.8$ with standard deviations never greater than 0.033. Thicker markets exhibit a superior allocative performance for lower values of γ but the overall conclusion is the same: a narrow (but not empty) price band is a necessary condition to attain sufficiently high levels of efficiency.

6 A Comparison of Alternative Rules

This section compares the performance of the protocol when adopting π -resampling versus γ -resampling over four different criteria: allocative efficiency (AE), mean spread, cancellation rate, and traders' protection (TP).

Table 3 compares the allocative efficiency under π - and γ -resampling for thin, thick, and crowded markets.

For each combination of n and resampling rule, we list the highest mean values obtained. These are slightly higher under γ -resampling, but we would not stake big claims over tiny differences that are subject to sampling errors. (However, they are statistically significant for $n = 200$ and $n = 1000$.) We prefer to conclude that there is no clear winner over AE : both resampling rules can be tuned to attain comparably high levels of allocative efficiency .

The second performance criterion is the mean spread. Figure 6 shows the best bid and ask under both π -resampling (on the left) and γ -resampling (on the right) with $n = 200$ agents. The y-axes are truncated to $[0.3, 0.7]$ to enhance readability.

Predictably, as a mere visual inspection confirms, the clear winner is the γ -resampling rule that is based on an explicit form of price control. Table 4 validates this conjecture by listing the lowest mean spread obtained under π - and γ -resampling for thin, thick, and crowded markets. The difference between the mean values is statistically significant for each choice of n .

Conditional on choosing the right parameter, the mean spread with γ -resampling is remarkably smaller. However, note that the best performances of both π - and γ -resampling with regard to the mean spread require a choice of parameters that are far from being optimal for allocative efficiency. This is easily seen by comparing the second rows from Tables 3 and 4. Therefore, while it is clear that γ -resampling yields a lower mean spread than π -resampling under ideal conditions, we need a further test to find out whether it is still superior once we take into account both efficiency and mean spread.




Fig. 6 Mean spread under π - and γ -resampling

Table 4 Mean spread under π - and γ -resampling

	π -resampling			γ -resampling		
	$n = 40$	$n = 200$	$n = 1000$	$n = 40$	$n = 200$	$n = 1000$
Minimum (mean) spread	0.267	0.152	0.113	0.228	0.118	0.086
Minimizer (π, γ)	0.000	0.000	0.000	0.450	0.500	0.600

This test is provided on the left-hand side of Fig. 7, where we plot the average outstanding bid and ask under both π -resampling (in red) and γ -resampling (in black) with $n = 200$ agents. This graph combines information about the two resampling rules. For each level of the (mean) allocative efficiency attained under either rule, we plot the corresponding average values of the best bid and ask and then join the datapoints using broken lines. Since the two rules attain different (mean) efficiencies, the datapoints are not vertically aligned. The left-hand picture shows clearly that, for comparable levels of allocative efficiency, γ -resampling leads to smaller (mean) spreads than π -resampling. In other words, a resampling rule based on a price band tends to produce a smaller spread than a rule based on a full clearing of the book, without sacrificing allocative efficiency.

Our third performance criterion is the cancellation rate. As is the case for π -resampling, this rate is on average increasing in γ because the width of the price band inversely correlates with the number of canceled orders. For $n = 200$, the mean (and standard deviation) of the cancellation rate go up from 1.868 (0.561) at $\gamma = 0.05$ (it is zero for $\gamma = 0$) to 18.16 (3.211) at $\gamma = 1$; the standard deviations are never greater than 3.212. A direct comparison shows that the range of attainable values for the cancellation rate is virtually identical under π - and γ -resampling. Similar results hold for $n = 40$ and $n = 1000$. Taken by itself, therefore, a criterion based on the cancellation rate is not conclusive.

As for the mean spread, however, we can compare the combined performance of either resampling rule with respect to allocative efficiency and cancellation rates.




Fig. 7 Bid-ask spreads and cancellation rates versus allocative efficiency

Table 5 Maximum traders' protection under π - and γ -resampling

	π -resampling			γ -resampling		
	$n = 40$	$n = 200$	$n = 1000$	$n = 40$	$n = 200$	$n = 1000$
Maximum TP	0.563	0.718	0.791	0.589	0.774	0.833
Maximizer (π, γ)	1.000	0.700	0.200	0.800	0.750	0.700

The right-hand side of Fig. 7 plots the (average) cancellation rates for each level of the (mean) allocative efficiency attained under either rule. For a large range of (lower) allocative efficiencies, π -resampling has a substantially lower cancellation rate; for high values, γ -resampling comes out better by a thin margin. (We do not report the graphs for different values of n , but increasing n makes this conclusion sharper.) Hence, whenever the market designer views the cancellation rate as ancillary to the allocative performance, he should prefer a resampling rule based on the price band.

The last (and in our opinion, more important) criterion is traders' protection. Table 5 compares the performance of π - and γ -resampling in thin, thick, and crowded markets. Similarly to Table 3, we list the highest mean values obtained for each combination of n and resampling rule. For $n = 200$ or $n = 1000$, the differences between the mean values are statistically significant. (For $n = 40$, this holds at the 1% significance level.) Conditional on choosing the right parameter, traders' protection is higher using γ -resampling. Note also that the optimal values of π and γ are decreasing in the thickness of the market, but this effect is much stronger for π -resampling. Therefore, when the exact size of the market is not known, the choice of the parameter under γ -resampling is more robust.

This superiority carries over when traders' protection is ancillary to allocative efficiency. The left-hand side of Fig. 8 superimposes the usual graphs of the mean values for AE and TP under γ -resampling for $n = 200$. The equivalent representation for π -resampling is on the right-hand side of Fig. 4. In general, γ -resampling




Fig. 8 Traders' protection and allocative efficiency for π - and γ -resampling

delivers a higher traders' protection than π -resampling for any given level of allocative efficiency. This is shown on the right-hand side of Fig. 8, where we report the (mean) traders' protection offered by the two resampling rules with respect to their (mean) allocative efficiency. The γ -resampling frontier on the AE - TP plane dominates the π -resampling frontier.

7 Conclusions

We have studied the continuous double auction from the point of view of market engineering, tweaking the trading protocol in search of improved designs. Our starting point has been the rules for exchange adopted by Gode and Sunder (1993a) for experiments with human agents and simulations with robot traders. We have disassembled their trading protocol into several component rules, and focused attention on resampling. We have assumed zero intelligence trading as a lower bound for more robust behavioral rules in order to elucidate the consequences of different resampling techniques.

Like Gode and Sunder (1993a) and most of the subsequent literature, we look first at allocative efficiency. Their trading protocol makes an extreme assumption that we call full resampling. We show that full resampling is especially favorable to allocative efficiency, biasing Gode and Sunder's results about the ability of the market to substitute for the lack of traders' intelligence. (A second negligible bias may come from their halting rule.) On the other hand, we demonstrate that partial resampling may be sufficient for the purpose of attaining a high allocative efficiency.

Based on this, we have devised a family of rules that includes as extreme cases both Gode and Sunder's full resampling and the opposite assumption of no resampling. This class of rules is parameterized by the probability π of clearing the book after a transaction occurs. We find that there is a large range of values for which π -resampling can attain a high allocative efficiency. In order to discriminate among

such π 's, we introduce three subordinate performance criteria: spread, cancellation rate, and traders' protection. The spread criterion measures the capacity of the protocol to provide a useful price signal. The cancellation rate looks at the inconvenience created by over-resampling. Finally, traders' protection measures the ability of a protocol to help agents capture their share of the competitive equilibrium profits. This latter criterion, patterned after the usual measure of allocative efficiency, is (to the best of our knowledge) new to the literature: we argue that ignoring it neglects one of the paramount goals of designing a market protocol.

We then introduce a different family of rules, based on the idea to delete only those quotes that fall outside of a price band parameterized by γ . We find that from the point of view of allocative efficiency, the optimized versions of either resampling rule are virtually indistinguishable. However, several differences emerge when we study their performance with respect to the other three criteria. In particular, when we consider a pair of criteria where the first one is allocative efficiency and the second one is any of the other three, we find that it is always the case that (at least for high efficiencies) γ -resampling dominates π -resampling. We then conclude that a resampling rule based on the price band is superior.

Acknowledgements This paper was written while the third author was visiting the School of Finance and Economics at the University of Technology of Sidney, whose financial assistance is gratefully acknowledged. We also received financial support from MIUR under grants 2007EE-NEAX and 2007TKLTSR.

References

- Arifovic, J., & Ledyard, J. (2007). Call market book information and efficiency. *Journal of Economic Dynamics and Control*, 31, 1971–2000.
- Brewer, P. J., Huang, M., Nelson, B., & Plott, C. R. (2002). On the behavioral foundations of the law of supply and demand: human convergence and robot randomness. *Experimental Economics*, 5, 179–208.
- Cliff, D., & Bruten, J. (1997). *More than zero intelligence needed for continuous double-auction trading*. Hewlett Packard Laboratories Paper HPL-97-157, Bristol, England.
- Crockett, S., Spear, S., & Sunder, S. (2008). Learning competitive equilibrium. *Journal of Mathematical Economics*, 44, 651–671.
- Duffy, J. (2006). Agent-based model and human subject experiments. In L. Tesfatsion & K. L. Judd (Eds.), *Handbook of computational economics* (pp. 950–1011). Amsterdam: Elsevier.
- Gjerstad, S., & Shachat, J. M. (2007). Individual rationality and market efficiency, wp 1204, August. IRBEMS, Krannert School, Purdue University.
- Gode, D. K., & Sunder, S. (1993a). Allocative efficiency of markets with zero intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101, 119–137.
- Gode, D. K., & Sunder, S. (1993b). Lower bounds for efficiency of surplus extraction in double auctions. In D. Friedman & J. Rust (Eds.), *The double auction market* (pp. 199–218). Reading: Addison-Wesley.
- Gode, D. K., & Sunder, S. (1997). What makes markets allocationally efficient? *Quarterly Journal of Economics*, 112, 603–630.
- Gode, D. K., & Sunder, S. (2004). Double auction dynamics: structural effects of non-binding price controls. *Journal of Economic Dynamics and Control*, 28, 1707–1731.

- Gul, F., & Lundholm, R. (1995). Endogenous timing and the clustering of agents' decisions. *Journal of Political Economy*, 103, 1039–1066.
- Hurwicz, L., Radner, R., & Reiter, S. (1975). A stochastic decentralized resource allocation process: part I. *Econometrica*, 43, 187–222.
- LiCalzi, M., & Pellizzari, P. (2008). Zero-intelligence trading without resampling. In K. Schrederseker & F. Hauser (Eds.), *Complexity and artificial markets* (pp. 3–14). Berlin: Springer.
- Maslov, S. (2000). Simple model of a limit order-driven market. *Physica A*, 278, 571–578.
- Mirowski, P. (2007). Markets come to bits: evolution, computation and markomata in economic science. *Journal of Economic Behavior and Organization*, 63, 209–242.
- Roth, A. E. (2002). The economist as engineer: game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70, 1341–1378.
- Smith, V. L. (1982). Microeconomic systems as an experimental science. *American Economic Review*, 72, 923–955.
- Stigler, G. J. (1964). Public regulation of the securities markets. *Journal of Business*, 37, 117–142.
- Subrahmanian, E., & Talukdar, S. N. (2004). Engineering of markets and artifacts. *Electronic Commerce Research and Applications*, 3, 369–380.
- Sunder, S. (2004). Markets as artifacts. In M. Augier & J. March (Eds.), *Models of a man: essays in memory of Herbert A. Simon* (pp. 501–520). Cambridge: MIT Press.
- Zhan, W., & Friedman, D. (2007). Markups in double auction markets. *Journal of Economic Dynamics and Control*, 31, 2984–3005.
- Zhan, W., Zhang, J., Yang, J., Wang, S., & Lai, K. K. (2002). *k*-ZI: A general zero-intelligence model in continuous double auction. *International Journal of Information Technology and Decision Making*, 1, 673–691.

Computational Methods in Economic Dynamics

Dawid, H.; Semmler, W. (Eds.)

2011, VIII, 216 p., Hardcover

ISBN: 978-3-642-16942-7