

Discussion Paper No. 1255

Strategy-proofness and Markets

by

Mark A. Satterthwaite

Northwestern University

March 10, 1999

<http://www.kellogg.nwu.edu/research/math>

Strategy-proofness and Markets*

Mark A. Satterthwaite
Kellogg Graduate School of Management
Northwestern University
Evanston, IL 60208

m-satterthwaite@nwu.edu
847-491-5482 (voice)
847-467-1777 (fax)

March 10, 1999

Abstract

If a market is considered to be a social choice function, then the domain of admissible preferences is restricted and standard social choice theorems do not apply. A substantial body of analysis, however, strongly supports the notion that attractive strategy-proof social choice functions do not exist in market settings. Yet price theory, which implicitly assumes the strategy-proofness of markets, performs quite well in describing many real markets. This paper resolves this paradox in two steps. First, given that a market is not strategy-proof, it should be modeled as a Bayesian game of incomplete information. Second, a double auction market, which is perhaps the simplest operationalization of supply and demand as a Bayesian game, is approximately strategy-proof even when the number of traders on each side of the market is quite moderate.

1 Introduction

Typically an intermediate microeconomics textbook defines the equilibrium price in a market as the price at which the supply and demand curves intersect, explaining that only at this price are both buyers and sellers satisfied: no buyer wants to offer a higher price in order to purchase an additional unit and no seller wants to offer a lower price in order to sell an additional unit. If the price in the market is not at the equilibrium level, then either some buyer or some seller will find it profitable to bid a shade higher or offer a shade lower with the

*Prepared as a keynote presentation for the Society for Social Choice and Welfare Conference held in Vancouver, British Columbia, in July 1998. I thank James Schummer for helpful discussions and Steven Williams for detailed and perceptive comments.

expectation that that change will permit it to trade additional units profitably.¹ Consequently only at the equilibrium price does the market come to rest.

This story of market equilibrium does not mention—in fact has no room for—strategic behavior. Supply and demand curves are built up from individual preferences on the assumption that each buyer (or seller) takes price as given and chooses quantity conditional on that price. Thus strategic behavior in which a buyer underreports the number of units he is willing to purchase at a given price in order to manipulate the price at which trade occurs is *a priori* ruled out. Implicitly the textbook story assumes that trade in a market is *strategy-proof*, i.e., each trader has a dominant strategy to report his or her preferences truthfully.

A market, however, is not generally strategy-proof as the following example illustrates. For the moment, let the market have two participants, a buyer and a seller. The seller has a single, indivisible widget on which she places a cost of $c \in [0, 1]$. The buyer would like to buy the widget if he does not have pay more than the value $v \in [0, 1]$ he places on it. The values c and v are private information to the seller and buyer respectively. Finally, let the equilibrium price be found through a simple well defined, auction-like process.

Suppose, for specificity, that $v = 0.8$ and $c = 0.55$. Since $v > c$ efficiency requires that trade take place. The auctioneer begins the process crying out a proposed price 0.5. The seller responds that she is unwilling to sell her widget at that price while the buyer responds that he is willing. Since excess demand exists in the market the auctioneer begins a process of raising the proposed price in increments of 0.1 until one of two events occurs: (i) the buyer and seller both accept the price and the market clears, or (ii) the buyer switches from accepting the proposed price p to rejecting it, in which case the auctioneer concludes that more than likely $v \not> c$, trade would be inefficient, and the market should close without trade occurring. Upon announcing the new proposed price of 0.6, the buyer responds that he is willing to buy. The seller, however, has a problem in deciding how to respond. She in fact is willing to trade at 0.6, but if she does agree to trade at this price, then her gain is only 0.05 (the price less her cost c).² She knows that the next price the auctioneer will cry out will be 0.7. Holding out for that price may be worthwhile, though it has the danger that the buyer will be unwilling to pay 0.7, an eventuality that would result in no trade because the auctioneer would incorrectly infer that $v < c$ and close the market. Nevertheless the seller takes the risk, the auctioneer announces 0.7, both the buyer and seller accept the price, and trade occurs with the seller realizing gain of 0.15 and the buyer realizing gain of 0.1.

This example illustrates two fundamental points: First, traders in a market may have an incentive to act strategically: the seller netted her an extra 0.1 gain at the buyer's expense as a result of her strategic maneuver. Second, strategic behavior may result in inefficiency. If the buyer's value had been 0.68 rather

¹See, for example, Frank (1991, p. 32–35).

²In the classical theory of tâtonnement as introduced by Walras she would accept the price 0.55, for that theory assumes price taking behavior on the part of all traders. See Negishi (1989, p. 589–96).

than 0.8, then the buyer would have surely rejected the price 0.7, the auctioneer would have closed the market, and the available gains from trade would have been lost as a result of the seller's maneuver.

Nothing about this example depends on the market consisting only of a single buyer and a single seller. Parallel examples are easily constructed for markets that have m buyers and m sellers—call such a market a size m market—in which each seller i has a one widget whose opportunity cost to her is $c_i \in [0, 1]$ and each buyer j seeks one widget whose value to him is $v_j \in [0, 1]$.

This example, showing the manipulability of a simple auction-like market, raises the following question: if market mechanisms that govern the setting of price are not strategy-proof, then why is supply and demand analysis (i.e., price theory) apparently so robust as an explanation for the prices that markets generate? The standard answer to this question, which text books introduce in their discussion of perfect competition, is that if the market is large in the sense of having many participants (none of whom is inordinately big), then no individual can significantly affect the price at which trade occurs. Therefore each trader is a price taker, has no incentive to misrepresent his or her preferences, and the market is strategy-proof. But this leaves open the question of how large does a market need be in order to be large, for strictly speaking an individual trader can affect price even in a very large finite market. It is this last question, how large is large enough, that is the focus of my discussion here. My belief is that the answer I propose provides insight into the possibilities of “approximate” strategy-proofness within market settings.

The argument that I make here depends on work that I have done with Aldo Rustichini and, especially, Steven Williams. It shows two things. First, if the text book concept of supply and demand is operationalized as a double auction under incomplete information within simple exchange markets of the sort used in the example above, then as the market size m increases the equilibrium strategic behavior in the market decreases rapidly towards zero. Specifically, if the market's size doubles, then the maximal amount by which any trader misrepresents his cost/value is cut in half. This convergence of strategies to truth-telling implies that each time the market size is increased by a factor of two, then in relative terms the expected inefficiency of the market decreases by a factor of four. Numerical examples suggest that an extremely good approximation to a large market may be obtained with a market of size $m = 8$. Second, in a well defined sense, the double auction does as well in achieving efficiency as any mechanism possibly can. In the simple exchange markets we have studied no mechanism converges to efficiency at a rate that dominates the quadratic rate the double auction achieves.

Before reviewing these results in detail, two topics need discussion. First, the assertion above that strategy-proofness is impossible in a market setting needs documentation. Second, the impossibility of “exact” strategy-proofness implies that the appropriate way to explore the possibility of “approximate” strategy-proofness is by modeling the market in question as a game of incomplete information and examining the nature of the equilibria. With these preliminaries complete, I review the convergence results we have obtained for double auctions

and then conclude with a brief discussion of some of the questions that these answers raise.

2 Strategy-Proofness within a Market

The standard, social choice formulation of strategy-proofness is this. There is a group of $m \geq 2$ individuals and a set $X = \{x, y, z, \dots\}$ of $n \geq 3$ alternatives among which they must choose one. Each individual i has a complete and transitive preference ordering R_i on X where, by notational convention, P_i is the strict, not necessarily complete ordering that R_i implies. Every ordering R_i of X is admissible; let Ω be the set of these possible orderings of X . The group uses the social choice function $\Psi : \Omega^m \rightarrow X$ to make its choice: Ψ maps profiles of reported preferences $R = (R_1, \dots, R_m) \in \Omega^m$ to a single alternative $x \in X$.

Define the notation (R_{-i}, \hat{R}_i) to refer to the profile R with, for individual i , the ordering \hat{R}_i substituted for R_i . The social choice function Ψ is *strategy-proof* if no profile $R \in \Omega^m$, no individual i , and no ordering $\hat{R}_i \in \Omega$ exists such that

$$\Psi(R_{-i}, \hat{R}_i) P_i \Psi(R),$$

i.e., no profile R exists at which i by reporting \hat{R}_i can force selection of the alternative $\Psi(R_{-i}, \hat{R}_i)$ that he prefers to $\Psi(R)$. If a social choice function is strategy-proof, then every individual's dominant strategy is to report his true preference ordering R_i . Strategy-proofness is attractive because it removes game theory—and its problematic predictions—from the group decision problem. Given this formulation the standard impossibility theorem for strategy-proof mechanisms is that no non-dictatorial, strategy-proof social choice function Ψ exists whose range includes every element of X .³

The standard theorem does not apply to markets because the market environment naturally imposes structure on preferences, so the set of admissible preferences for an individual is some proper subset of Ω . For example, consider a size m market with m sellers and m buyers in which the traders are assumed to have transferrable utility. Specifically, the seller's utility is

$$u_S(c_i, \delta_i, t_i) = s_i - (1 - \delta_i)c_i \tag{1}$$

where $c_i \in [0, 1]$ is the opportunity cost she incurs in selling the widget, $\delta_i \in \{0, 1\}$ is an indicator variable that takes on the value 1 if she does not sell her widget (i.e., keeps it) and 0 if she sells it, and $s_i \in \mathfrak{R}$ is the monetary transfer she receives. Thus if she sells her widget at price p , then her utility is $p - c_i$ while if she does not sell her widget and neither makes nor receives a transfer, her utility is 0. Similarly, a buyer's utility is

$$u_B(v_j, \lambda_j, t_j) = \lambda_j v_j + t_j \tag{2}$$

³See Gibbard (1973) and Satterthwaite (1975). A social choice function F is dictatorial if an individual i (the dictator) it always picks as the social choice an alternative that is maximal relative to i 's reported preferences R_i .

where $v_j \in [0, 1]$ is the value he realizes in buying the widget, $\lambda_j \in \{0, 1\}$ is an indicator variable that takes on the value 1 if he purchases the widget and 0 if he fails to purchase it, and $t_j \in \mathfrak{R}$ is the monetary transfer he receives.⁴ Note that these utility functions have only c_i and v_j as parameters; consequently c_i and v_j fully describe the seller and buyer's preferences respectively.

Feasibility implies that the set of alternatives among which the social choice function selects is

$$X^{**} = \{(\delta_1, \dots, \lambda_m, s_1, \dots, t_m) \mid \Sigma_i \delta_i + \Sigma_j \lambda_j \leq m, \Sigma_i s_i + \Sigma_j t_j \leq 0\}. \quad (3)$$

where the constraints follow from three facts: only m widgets are available for trade, no outside source of funds is available to subsidize the market, and widgets and money are freely disposable. X^{**} obviously has an infinite number of elements; thus Ω , the set of all possible orderings on this set X^{**} of alternatives, is inconceivably large. This contrasts with the situation here: given that c_i and v_j are *a priori* restricted to the unit interval and parameterize the admissible utility functions $u_S(c_i, \cdot)$ and $u_B(v_j, \cdot)$, the sets Ω_S^{**} and Ω_B^{**} (the admissible preferences for the seller and buyer respectively) are negligible subsets of Ω .

The importance of this is that if the set of admissible preferences are small compared to Ω , then the possibilities for constructing a strategy-proof social choice function are enhanced because there are a limited number of orderings $\hat{R}_i \in \Omega_S^{**}$ available for a seller to choose among in order to manipulate the outcome in her favor. Barberà and Jackson (1995) have written a remarkable paper characterizing the strategy-proof social choice functions that exist for an exchange economy: each of the m agents receives an endowment of the l goods that exist in the economy, each agent has a strictly quasi-concave, continuous, and increasing utility function u_i on \mathfrak{R}^l (denote this class of utility functions as Ω^*), and the social choice function prescribes a set of net trades among the agents as a function of their reported utility functions. Their main result (Theorem 3) shows that limiting admissible preferences to Ω^* does succeed in expanding the class of strategy-proof social choice functions beyond dictatorial functions.⁵ Specifically, the theorem establishes that, in this market setting, an anonymous, non-bossy social choice function is strategy-proof if and only

⁴Normally $t_j < 0$ if $\lambda_j = 1$, for if the buyer succeeds in buying the widget, then $-t_j$ is the price he pays for it.

⁵Two comments are merited concerning dictatorial social choice functions. First, a dictatorial social choice function is not anonymous because it identifies a particular individual as the the dictator. Second, within a market setting a dictatorial social choice function can be quite complex because agent i only cares about his own allocation, not the allocations of other agents. This restriction permits the construction of strategy-proof social choice functions that are *serially dictatorial*. In a serially dictatorial social choice function a fixed heirarchy of agents exists. The first agent at the top of the heirarchy (the top level dictator) chooses his most preferred allocation from a set that is exogenously fixed as part of the functions's definition. The next agent down the heirarchy (the second level dictator) chooses his most preferred allocation from a set that is a function top level dictator's choice, etc. Such functions are strategy-proof because no pair of agents mutually affects the set of alternatives from which each chooses; therefore no strategic interaction exists.

if it prescribes anonymous fixed proportion trading.⁶ A social choice function is anonymous if the net trades prescribed to an agent do not depend on the identity of that individual. It is non-bossy only if an agent must change his own allocation in order to affect another agent's allocation.

What is anonymous fixed proportion trading and, given that, what do we really gain from limiting admissible preferences from Ω , the set of all possible orderings of \mathcal{R}^I , to Ω^* , the set of strictly quasi-concave, continuous, and increasing utility functions on \mathcal{R}^I ? Fixed proportion trading is complicated to define in general; consequently I demonstrate it here within the simple supply-demand example used above using four traders (two sellers and two buyers) even though the only if part of Barberà and Jackson's theorem does not apply to this special case.⁷

The two sellers have costs $c_1, c_2 \in [0, 1]$ and the two buyers have values $v_1, v_2 \in [0, 1]$. Efficiency requires that the two widgets, initially owned by the two sellers, be assigned to the two traders who have the highest values. Thus if the preference profile is $c_1 = 0.3$, $c_2 = 0.8$, $v_1 = 0.6$, and $v_2 = 0.4$, as is shown in Example 1 of Figure 1, then efficiency requires buyer 1 and seller 2 be assigned the two widgets, i.e., efficiency requires buyer 1 and seller 1 to trade. The simplest member of the fixed proportions class of trading rules is the fixed price rule.⁸ Given an *a priori* chosen, fixed price p_F , the rule prescribes that every buyer who reports a valuation greater than p_F trades, every seller who reports a cost less than p_F trades, and if the market demand and supply does not balance at p_F , then the long side of the market is rationed by random draw.⁹ It is strategy-proof because no agent can affect the terms of trade. Let, for the examples that follow, $p_F = 0.5$. In Example 1 the fixed price rule achieves efficiency, for seller 1 does sell to buyer 1 as efficiency requires.

The fixed proportions rule requires that three prices be given a priori: $p_{B3} > p_2 > p_{S3}$. The procedure for choosing among these prices is:

1. If exactly 3 traders (whether buyers or sellers is irrelevant) have valuations/costs greater than p_{B3} , then the price is fixed at p_{B3} , and trade—

⁶The class of social choice functions and environments that they consider is restricted to include only functions and environments that satisfy two technical conditions: tie-freeness and a restriction on the inequality of endowments.

⁷ Ω_S^* and Ω_B^* , the sets of preference orderings this example admits is drastically restricted from Ω^* . Therefore Barberà and Jackson's characterization may be incomplete for this set, and in fact it is. McAfee (1992) has devised a clever strategy-proof mechanism that uses dual prices, runs a surplus, and is not identified by Barberà and Jackson's theorem. Williams (1998), however, has shown that this restriction of preferences is not sufficient to permit the construction of an efficient and strategy-proof mechanism.

⁸Hagerty and Rogerson (1987) introduced the fixed price rule to the strategy-proofness literature.

⁹For example, suppose $p_F = 0.5$ and $c_1 = 0.3$, $c_2 = 0.8$, $v_1 = 0.7$, and $v_2 = 0.6$. Buyers 1 and 2 both want to trade at this price, but on the other side only seller 1 is willing. The choice between buyers 1 and 2 is random: with probability 0.5 buyer 1 trades and with probability 0.5 buyer 2 trades. Note that this random rationing is inefficient: buyer 1 should receive priority because his value is greater than 2's value. If the rationing were allowed to depend on the reported valuations, then the rule would no longer be strategy-proof.




Figure 1: Three examples of fixed proportions trading compared with fixed price trading.

including necessary rationing—is conducted at that price exactly as in the fixed price rule.

2. If exactly 2 traders have valuations/costs greater than p_2 , then the price is fixed at p_2 , and trade is conducted at that price exactly as in the fixed price rule. Or, equivalently, if exactly 2 traders have valuations/costs less than p_2 , then the price is fixed at p_2 , and trade is conducted at that price exactly as in the fixed price rule.
3. If exactly 3 traders have valuations/costs less than p_{S3} , then the price is fixed at p_{S3} , and trade is conducted at that price exactly as in the fixed price rule.
4. If none of the above apply, then no trade occurs.

This rule is strategy-proof because any trader who successfully shifts the price in his favor through misrepresentation of his true preferences necessarily causes

himself to be excluded from trading. To demonstrate the rule, let the three prices p_{B3} , p_2 , and p_{S3} be, as shown in Figure 1, 0.75, 0.50, and 0.25 respectively. Applying these rules to Example 1 results in the price $p_2 = 0.5$ being fixed because exactly two traders have valuations greater than p_2 . The efficient allocation then ensues exactly as with the fixed price rule.

Examples 2 and 3 show how the fixed proportions rule differs from the fixed price rule. In Example 2 the fixed proportions rule prescribes no trade because only 2 traders have valuations/costs greater than p_{B3} while 3 traders have valuations greater than p_2 . The fixed price rule does better: seller one sells his widget to either buyer 1 with probability 0.5 or buyer 2 with probability 0.5. Neither rule, however is efficient; that requires buyer 1 to trade with seller 1. The situation reverses in Example 3. There the fixed proportions rule selects price p_{B3} because 3 traders have valuations greater than p_{B3} . This results in seller 1 selling her widget and buyers 1 and 2 being rationed randomly. This rationing is not efficient, for buyer 1, with his value being greater than buyer 2's value, should trade with certainty.

These two examples illustrate that neither rule is efficient and, neither rule dominates the other. With respect to the first point it has long been known that severe difficulties exist in obtaining both strategy-proofness and efficiency in market settings. Hurwicz (1972) devised the initial demonstration of the incompatibility of the two requirements for a two good-two agent model of an exchange economy. Hurwicz and Walker (1990) show the impossibility of strategy-proof, efficient mechanisms for exchange economies in which agents are restricted to have quasi-linear preferences. Barberà and Jackson (1995), as we have been discussing, characterize for exchange economies anonymous, non-bossy, and strategy-proof social choice functions over a very general class of exchange economies and observe that the resulting rules are not efficient. Schummer (1997) has recently shown that for the case of two agents even the extreme step of restricting preferences to be linear is not sufficient to obtain both strategy-proofness and efficiency. Therefore, even though a completely general theorem has not yet been formulated and proved, it seems clear that no attractive social choice functions exist for market settings that are both strategy-proof and efficient.

With respect to the second point, Examples 2 and 3 show that neither the fixed price nor the fixed proportions rule is more efficient than the other for all possible preference profiles. A simple tabulation of the possibilities reveals that situations such as Example 2, in which the fixed proportions rule leads to no-trade and the fixed price rule realizes some of the available gains from trade, are much more prevalent than situations such as Example 3, in which the fixed proportions rule realizes gains from trade and the fixed price rule leads to no-trade.¹⁰ This observation suggests that the fixed proportions rule, despite the apparent flexibility it imparts on price, may in fact perform more poorly

¹⁰The reason is that the event "exactly three traders reporting costs/values greater than p_{B3} or exactly two traders reporting values/costs greater than p_2 or exactly three traders reporting values/costs less than p_{S3} " is quite a restrictive event and can easily fail to be realized even when trade should occur.

in terms of efficiency than the fixed price rule with its single price. Thus our deep and extremely ingenious explorations of the possibilities for constructing strategy-proof social choice functions in market environments have brought little reward in terms of discovering usable mechanisms.¹¹

Return now to the question that I posed at the beginning: Why is price theory, which implicitly assumes truthful revelation of preferences, such a robust explanation for the prices that we observe in markets? This short review has shown that it is not and can not be that market mechanisms are strategy-proof. Yet, as price theory's success suggests, people often act as if markets were strategy-proof. How can this paradox be resolved?

3 Strategy-Proofness, Games of Incomplete Information, and Mechanism Design

Given that market mechanisms are not and can not be strategy-proof and given that price theory does a pretty fair job of explaining prices without any reference to traders' efforts to manipulate the outcome, then how should the trading process be modeled so that the apparent unimportance of strategic behavior in the price setting process be investigated rather than assumed. My starting point is the observation that if a social choice function Ψ is not strategy-proof, as is normally the case in a market setting, then each individual's best response to other agents is a function of their preferences. This follows directly from the definition of manipulability: given that the mechanism Ψ is not strategy-proof, then an individual i , profile $R \in \Omega^m$, and ordering $\hat{R}_i \in \Omega$ exist such that

$$\Psi(R_{-i}, \hat{R}_i) P_i \Psi(R), \quad (4)$$

which is to say that if $\sigma_i : \Omega^m \rightarrow \Omega$ is i 's best response correspondence, then $\sigma_i(R_{-i}|R_i) \neq R_i$ but rather $\sigma_i(R_{-i}|R_i) = \hat{R}_i$.

This almost trivial observation has a deeper implication: each individual has good reason to be uncertain about other individual's preferences. To see this, let ψ be the manipulable social choice function a group is using and suppose that whatever preference profile $R \in \Omega^m$ is realized all m individuals believe they know every other individual's preferences. Given a realized preference profile $R \in \Omega$, the reported preference profile is the complete information Nash equilibrium, $\Sigma(R)$.¹² Therefore the composition of ψ and Σ gives the outcome; let this function be $\Psi(R) \equiv \psi(\Sigma(R))$. The function Ψ is itself a social choice

¹¹I use the word "our" literally, for I participated in one of the earlier efforts. See Satterthwaite and Sonnenschein (1981).

¹²I am assuming here that the Nash equilibrium is in pure strategies and that, if multiple equilibria exist, Σ makes a selection among them. If the function f were one for which, for some profile R only mixed strategy equilibria exist, then this argument would have to be recast using a characterization of strategy-proofness for social choice functions that specify lotteries as well as singletons. Gibbard (1978), for example, characterized such functions for the standard social choice setting of a discrete set of alternative with no *a priori* structure on the admissible orderings.

function and, unless ψ and Σ have a very fortuitous structure, Ψ is not strategy-proof. But if Ψ is not strategy-proof, then an individual i , a preference profile R , and ordering \hat{R}_i exists such that (4) holds. The interpretation of (4) here is that if agent i anticipates that profile R is likely to be realized and if he can mislead everyone else into believing that his preferences are \hat{R}_i and not R_i , then he can secure the outcome $\Psi(R_{-i}, \hat{R}_i)$, not the outcome $\Psi(R)$ that he prefers less according to his true preferences R_i .¹³ In other words, agent i manipulates the social choice by manipulating others' beliefs about his preferences.

This possibility of manipulating other individual actions by misleading them in the formation of their beliefs suggests two complementary observations.

- This type of manipulation is common in the everyday world. Consider three examples. A frequent, and sometimes effective way of bargaining is for the buyer to convince the seller that he does not care very strongly for the good that is being offered in the expectation that she will reduce the asking price in order to clinch the deal. The essence of a confidence man's swindle is to convince the victim absolutely that he has her best interests at heart. Similarly, an assistant professor may seek to convince his senior colleagues that he has a more genuine love for scholarship than in fact he does, for that belief may make his colleagues more willing to grant tenure with its commitment of lifetime employment.
- Most people have been manipulated in exactly this way. Few of us are not gullible on occasion. A reasonable response to such experiences is to become more cautious in believing that one really knows what another person's preferences are. As a consequence, I would assert, experienced and rational people tend to retain a degree of scepticism about the preferences of even people they know quite well. This is especially true for a person in a position of authority, for she certainly realizes that the people over whom she has some measure of power tend to tell her what she wants to hear.

The implication of these observations for understanding the robustness of price theory is straightforward: any adequate theory must explicitly take into account the uncertainty traders have concerning other traders' preferences and explain how the market mechanism successfully extracts this preference information even as it uses this information to allocate the widgets being traded.¹⁴ The essence of the exchange problem is asymmetric information about preferences.

The obvious way to proceed, which I follow, is to model trade as a Bayesian game of incomplete information and to derive its equilibrium properties. This, however, is not the only way to proceed. For example, Jackson and Manelli (1997) use less structure on beliefs than the Bayesian game approach requires and show how an exchange economy must converge to a competitive equilibrium

¹³This argument to my knowledge was first developed in Blin and Satterthwaite (1975).

¹⁴This rules out the theory of Nash implementation as a useful tool except in special situations.

as the number of agents become large. The generality of their results, however, exacts a price: they do not derive any rates at which equilibrium allocations converge to an efficient allocation as the economy grows. Therefore they can not give an indication as to how big a market must be in order to be approximately competitive.

Before turning to a discussion of the results that have been achieved in the Bayesian game framework, one additional modeling issue should be discussed. Mechanism design theory is a marvelous tool that has been developed over the past 20 years that allows us to understand a host of issues involving asymmetric information. It, however, is not directly useful here, for it assumes that the mechanism designer as well as the participants share common knowledge of the underlying distribution F of preferences. Taking the prior distribution F as given, the mechanism designer might construct, for example, a social choice function $\Psi : (\Omega_S^{**})^m \times (\Omega_B^{**})^m \rightarrow X^{**}$ such that it maximizes the ex ante expected gains from trade subject to the constraints of:

- Incentive compatibility, which requires that each trader's Bayesian Nash equilibrium strategy be to report his true cost/value,
- Interim individual rationality, which requires that each trader's expected utility conditional on his cost/value be nonnegative, and
- Ex ante budget balancing, which requires that the expectation of the summed monetary transfers $\sum_i t_i + \sum_j t_j$ be nonpositive.

As long as the underlying prior distribution F remains constant Ψ remains constrained efficient and this works well. If, however, the underlying prior should change from F to \hat{F} , then one of two things happens:

- The social choice function Ψ remains fixed, and as a consequence of the distribution changing from F to \hat{F} , each agent changes his strategy and honest revelation of preferences is no longer necessarily a Bayesian Nash equilibrium. There is no reason to think that the resulting equilibrium achieves, or comes close to achieving, the same constrained efficiency as the original equilibrium did for the original distribution F .
- The mechanism designer reoptimizes the social choice function from Ψ to $\hat{\Psi}$ based on his understanding of what the new distribution is. If the new prior \hat{F} is common knowledge among everyone, agents and mechanism designer alike, then the redesign is successful and there is nothing more to be said. If, however, the mechanism designer does not share in the common knowledge among the traders that the distribution is now \hat{F} , then it is unclear how he can possibly obtain the needed information from the self-interested agents. The mechanism Ψ is only designed to elicit preferences information given common knowledge of the underlying distribution, not to elicit sequentially information about the agents' beliefs about each other and then about their individual, private preferences.

An appropriate approach to avoid these conundrums is to abandon the mechanism design approach of tailoring the mechanism to a specific environment and instead to pick Ψ so that, in some sense, it performs well no matter what the distribution of preferences turns out to be. This is the approach that Wilson (1987, p. 36) has advocated, observing that the rules of a market “are not changed as the environment changes; rather they persist as stable, viable institutions.” I now turn to showing that the double auction fulfills this goal. While neither strategy-proof nor efficient, it is approximately strategy-proof even for moderately sized markets and approaches full *ex post* efficiency with surprising speed as the market grows in size.

4 Properties of the Double Auction

4.1 Model

The double auction is a simple operationalization of supply and demand as a one-shot game that Aldo Rustichini, Steven Williams, and I have studied. There are m sellers, each of whom owns a single, indivisible widget, and m buyers, each of whom might like to purchase a single widget. For a given seller i the widget’s opportunity cost, which must be met if she is to be willing to transfer it, is c_i . This cost is private to her; the other $2m - 1$ traders believe that her cost c_i is independently drawn from the distribution $F(\cdot)$ on the unit interval $[0, 1]$. For a given buyer j the widget’s value is v_j ; the other $2m - 1$ traders believe that it was drawn independently from the distribution $G(\cdot)$ on $[0, 1]$. We assume that F and G are C^1 and that their densities f and g respectively are bounded above by the positive constant \bar{q} and below by the positive constant \underline{q} . The pair (F, G) , which we call the market’s *environment*, along with the rules of the double auction and the strategies of each trader are common knowledge.

The sequence of five events that constitute a double auction are:

1. The sellers and buyers all independently draw their costs c_i and values v_j . Each trader’s cost/value is private information to that trader.
2. Simultaneously sellers and buyers submit offers and bids that are functions of their costs/values. Specifically, $S_i(c_i)$ is seller i ’s offer when her cost is c_i and the function $S_i : [0, 1] \rightarrow [0, 1]$ is her strategy. Exactly parallel, $B_j(v_j)$ is buyer j ’s bid when his value is v_j and the function $B_j : [0, 1] \rightarrow [0, 1]$ is his strategy.
3. Given the offers and bids a market clearing price is set. Sort the $2m$ offers/bids from lowest to highest,

$$s_{(1)} \leq s_{(2)} \leq \dots \leq s_{(m)} \leq s_{(m+1)} \leq \dots \leq s_{(2m)},$$

and fix the price at

$$p = \frac{s_{(m)} + s_{(m+1)}}{2}, \tag{5}$$

which is market clearing with respect to traders' reported costs and values.¹⁵

4. Assign the m widgets to the m traders that made the m highest offers/bids. Each buyer who acquires a widget pays p and each seller who relinquishes a widget receives p . Consistent with the utility functions (1) and (2) a seller who sells her widget realizes utility $p - c_i$, and a buyer who purchases a widget realizes utility $v_j - p$. Sellers and buyers who do not succeed in trading realize 0 utility. Note that summing the traders' utilities gives the total gains from trade that the double auction realizes.
5. The market disperses. Traders do not have another chance to either purchase or sell widgets.

The one-shot nature of the game is crucial, for if it were repeated, then each trader's opportunity cost of trading would not be c_i or v_j , but would be a transformation of c_i or v_j reflecting future opportunities to trade, perhaps at a better price.

Modeled in this manner the double auction is a game of incomplete information and the appropriate solution concept is Bayesian Nash equilibrium. A set of strategies, $\langle S_1, \dots, S_m, B_1, \dots, B_m \rangle$ is an equilibrium if (i), for each seller i and for each possible cost $c_i \in [0, 1]$, the offer $S_i(c_i)$ is a best response to the strategies of the $2m - 1$ other traders and (ii), for each buyer j and for each possible value $v_j \in [0, 1]$, the bid $B_j(v_j)$ is a best response to the strategies of the $2m - 1$ other traders. From the viewpoint of an individual trader the bids and offers of the other traders are random variables because their costs and values are private information. Therefore, given the strategies of the other traders, an offer $S_i(c_i)$ is a best response if it maximizes the expected utility of seller i when her cost is c_i . Similarly a bid $B_j(v_j)$ is a best response if it maximizes buyer j 's expected utility when his value is v_j .

We only consider symmetric equilibria in which all sellers play the same strategy $S(\cdot)$ and all buyers play the same strategy $B(\cdot)$. We also only consider equilibria in which trade occurs with positive probability and equilibria in which sellers always ask at least as much as their cost and buyers offer no more than their value.¹⁶ This structure implies that in equilibrium (i) sellers and buyers' strategies are strictly increasing over the relevant domains and (ii) a tie between a bid and an offer is a zero probability event.

¹⁵More generally price can be set anywhere in the interval $(s_{(m)}, s_{(n+1)})$ of possible market clearing prices according to the formula, $p = (1 - k)s_{(m)} + k s_{(m+1)}$, where $k \in (0, 1)$ is a fixed parameter that is common knowledge.

¹⁶The first part of this statement rules out the uninteresting no-trade equilibrium in which each seller, no matter what her cost is, asks for a price 1 and each buyer, no matter what his value is, demands a price 0. The second part of the statement rules out dominated strategies in which, for instance, a seller with cost c'_i that is near the right end of the interval $[0, 1]$ offers to sell her widget for less than her cost ($S(c'_i) < c'_i$). Her equilibrium strategy S might validly specify such an offer because in equilibrium the highest bid any buyer will make is $B(1)$, which will be strictly less than 1, and if $S(c'_i) > B(1)$, then her offer will never be accepted, i.e., the dominated offer $S(c'_i)$ is costless for her to make.

Given an environment (F, G) and a selected equilibrium $\langle S, B \rangle$ for a size m market, the double auction's *relative inefficiency* is

$$e(\phi_m^{\text{DA}}, F, G) = \frac{\Gamma_m(F, G) - \phi_m^{\text{DA}}(F, G)}{\Gamma_m(F, G)}$$

where (i) $\Gamma_m(F, G)$ is the ex ante expected gains from trade if each trader acted as if the double auction were strategy-proof and reported their true costs/values, thus guaranteeing an *ex post* efficient outcome in which all potential gains from trade are always realized, and (ii) $\phi_m^{\text{DA}}(F, G)$ is the ex ante expected gains from trade that the selected equilibrium achieves. In short, $e(\phi_m^{\text{DA}}, F, G)$ is the proportion of the possible gains from trade that the double auction's selected equilibrium $\langle S, B \rangle$ fails to achieve in expectation.

4.2 Approximate Strategy-Proofness of the Double Auction

The main theorem that Rustichini, Williams and I proved concerning equilibria of double auction is this.¹⁷

Theorem 1 (*Rustichini, Satterthwaite, and Williams, 1994*). *Given an environment (F, G) , continuous, positive functions $\kappa(q, \bar{q})$ and $\xi(q, \bar{q})$ exist such that, for all market sizes $m \geq 2$ and all equilibria $\langle S, B \rangle$ of the double auction,*

$$S(c_i) - c_i \leq \frac{\kappa(q, \bar{q})}{m} \tag{6}$$

for all $c_i \in [0, \bar{c}]$,

$$v_i - B(v_i) \leq \frac{\kappa(q, \bar{q})}{m} \tag{7}$$

for all $v_i \in (\underline{v}, 1]$, and

$$e(\phi_m^{\text{DA}}, F, G) \leq \frac{\xi(q, \bar{q})}{m^2}. \tag{8}$$

Inequalities (6) and (7) say that in equilibrium the maximal amount by which traders misrepresent their costs/values—i.e., misrepresent their preferences—is cut in half each time the market's size doubles. Note that these bounds apply to all possible equilibria, even equilibria with discontinuous strategies if such equilibria should exist. Inequality (8) says that the maximal relative inefficiency of equilibria is cut by a factor of four each time the market's size doubles.

¹⁷In their statement of the theorem the bounds (q, \bar{q}) on the densities f and g do not appear. They are added here because they are necessary in the statement of Theorem 2 below. It is a tedious but straightforward exercise to go through the proofs and derive how the constants κ and ξ depend on (q, \bar{q}) .

The driving force behind the bounds (6) and (7) on strategic misrepresentation can be explained as follows. Taking the viewpoint of a buyer, his gain from misrepresenting his value v_i downward is to try to force the market clearing price p downward. But, given the rule (5) for computing p , he can only do this if his bid is the $(m + 1)$ st highest among the $2m$ bids and offers.¹⁸ The probability of being the $(m + 1)$ st highest offer/bid is roughly proportional to $\frac{1}{m}$, i.e., it decreases quickly with market size. On the other hand the cost of misrepresenting is that he may be inefficiently excluded from trading, which occurs if $B(v_i) < p < v_i$. For a given magnitude of misrepresentation, $v_i - B(v_i)$, the probability of being excluded does not necessarily decrease with market size, but may actually increase because the standard deviation of p around the limit competitive price is proportional to $\frac{1}{\sqrt{m}}$. The bounds (6) and (7) in effect represent a balance, as m increases, between the decreasing expected gains from misrepresentation and the costs of misrepresentation.

Figure 2 shows for the uniform environment (i.e., F and G are both the uniform distribution) the effect of these bounds by graphing a representative bundle of equilibrium strategies for markets of sizes 2, 4, and 8 respectively.¹⁹ The horizontal axis represents traders' costs/values and the vertical axis represents their offers/bids. The 45° diagonal line represents truth-telling; if the double auction were strategy-proof, then that line would represent every trader's strategy. Because traders do misrepresent, sellers' strategies are above the diagonal and buyers' strategies are below the diagonal. Comparison of the three panels for different market sizes shows two effects of increasing market size: (i) traders' strategies approach truth-telling and (ii) the bundles of strategies rapidly become smaller. Point (ii) is interesting because it suggests that in an important sense as m becomes large the double auction has an essentially unique equilibrium.

Figure 3—a Harberger triangle—supplies intuition as to why the linear bounds (6) and (7) on misrepresentation lead to the quadratic bound (8) on relative inefficiency. Let the environment be the uniform environment. The SS and DD curves represent the expected number of widgets sold per seller and purchased per buyer, i.e., SS and DD are the expected per capita supply and demand curves for a double auction in the uniform environment. If the double auction were strategy-proof and if traders reported their true costs/values, then at point A the expected market price would be $p = 0.5$, the expected per capita quantity would be $q = 0.5$, and the allocation would be ex post efficient.

But traders do misrepresent. For market size m' a seller's offer is less than 0.5 only if $c_i < c'$ because $S(c') = 0.5$ and, similarly, a buyer's bid is greater than 0.5 only if $v_j > v'$ because $B(v') = 0.5$. Therefore the expected per capita quantity sold is not 0.5, but rather is q' because (i) all sellers for whom $c_i \in (c', 0.5)$ should trade but do not and (ii), similarly, all buyers for whom $v_j \in (0.5, v')$ should trade but do not. The area of the triangle ABC therefore represents the

¹⁸A buyer may affect the price if his bid is the m th highest, but in that case it does him no good because he does not trade.

¹⁹The best that can be done is to plot a representative bundle of equilibria because numerical experimentation suggests that a double continuum of smooth equilibria exist.

expected per trader gains from trade that are not realized as a result of strategic behavior.

Double the market size to $2m'$ and let the new equilibrium strategies be \tilde{S} and \tilde{B} . The per capita expected supply and demand curves remain unchanged. Let $\tilde{S}(c'') = 0.5$ and $\tilde{B}(v'') = 0.5$ where

$$0.5 - c'' \approx \frac{0.5 - c'}{2} \text{ and } v'' - 0.5 \approx \frac{v' - 0.5}{2}$$

because the magnitude of each traders' misrepresentation is cut in half approximately. The expected numbers of units sold per capita is therefore q'' where $0.5 - q'' \approx \frac{0.5 - q'}{2}$. In conformance with the theorem's bound (8) on the relative inefficiency of the double auction, the area of triangle AEF, which represents the expected per trader gains from trade that are not realized as a result of strategic behavior, is one-fourth the area of triangle ABC.

The power of this bound is shown in Table 1. For the case of the uniform environment it shows the relative inefficiencies of three different allocation mechanisms:

- The constrained efficient mechanism in which mechanism design theory has been used to tailor the mechanism to minimize its relative inefficiency subject to three constraints: incentive compatibility, interim individual rationality, and ex ante budget balancing.
- The double auction. The two values shown for its relative inefficiency are the minimum and maximum inefficiencies found within a representative bundle of equilibria.
- The fixed price rule with the fixed price set at the competitive price, $p = 0.5$, of the limit market with $m = \infty$.

Inspection of the table shows that for m as small as 6 or 8 both the constrained efficient mechanism and the double auction are essentially ex post efficient. Consequently strategic behavior within the double auction is not a practical issue for markets even as small as that.

Table 1: Relative inefficiencies of three mechanisms

m	Constrained Efficient Mechanism	Double Auction		Fixed Price Rule
		Most Efficient	Least Efficient	
2	0.056	0.056	0.063	0.22
4	0.015	0.015	0.016	0.18
6	0.0069	0.0069	0.0070	0.16
8	0.0039	0.0039	0.0039	0.15

This last observation, that for m on the order of 6 or 8 strategic behavior is inconsequential, is the answer that I offer to the question I originally posed:

why is price theory so robust even though it ignores the possibility of strategic behavior? This model of double auction mediated trade suggests that the effects of strategic behavior attenuate so quickly as the size of the market increases that it is a good and justifiable approximation to ignore its existence even in moderate sized markets.

The double auction achieves this excellent performance by being *approximately strategy-proof* in moderate and larger sized markets. Without offering a formal definition, I suggest that it is approximately strategy proof in two senses. First, as m increases equilibrium strategies converge towards truth-telling at a fast rate. Thus, for moderate sized and greater m , equilibria are almost strategy-proof in the sense that the equilibrium behavior of every trader is to reveal almost his true preferences. Second, if traders are boundedly rational (or just lazy), then substituting the focal strategy of truthful revelation for the equilibrium strategies is a reasonable decision since, for all but genuinely small markets, truthful revelation is an ε -equilibrium that, additionally, happens to be ex post efficient. Moreover it is an ε -equilibrium that is immune to the difficulty that Gul and Postlewaite (1991, p. 1284) raise concerning such ε -equilibria:

. . . while approximately optimal behavior may result in efficient outcomes, this certainly does not imply that precisely optimal behavior will result in approximately optimal outcomes. The cumulative effect of many agents' adjustments from approximately optimal behavior to optimal behavior, the subsequent adjustments to these adjustments, etc., can be large.

For the double auction we know that these adjustments are small and that precisely optimal behavior leads to an equilibrium that is close to the truthful revelation outcome. Thus even if some traders precisely optimize and others treat the double auction as strategy-proof and truthfully reveal, the outcome remains quite satisfactory from a welfare perspective.

4.3 Worst-Case Asymptotic Optimality of the Double Auction

The above discussion establishes that the double auction performs exceedingly well. Nevertheless if one is to paste the label "approximately strategy-proof" onto it for moderately sized and larger markets, then one would like to know that there is not some other allocation mechanism Ψ that is usable across a variety of environment (F, G) and that dominates the double auction in the quickness of the convergence of its equilibria to ex post efficiency. In recent work Steven Williams and I (1998) have shown that in a well defined sense this is not a issue: an environment (F, G) exists in which the double auction's rate of convergence to efficiency is equal to the rate at which the best possible allocation mechanism asymptotically converges to efficiency.

The framework for this result is as follows. Let the set of admissible environments (F, G) be E . A *market game* ϕ_m for a size m market consists of a

strategy set for each of the m sellers and m buyers, an outcome function that maps strategy profiles into allocations, and an equilibrium selection rule that for each environment in E picks a Bayesian Nash equilibrium of the market game. A *market mechanism* Φ over E is a sequence of market games for each possible market size $m \geq 2$; thus $\Phi = \{\phi_2, \phi_3, \dots\}$. Let $\phi_m(F, G)$ be the ex ante expected gains from trade that the market game ϕ_m generates in the environment (F, G) and, as before, let the mechanism's relative inefficiency be

$$e(\phi_m, G, F) = \frac{\Gamma_m(G, F) - \phi_m(F, G)}{\Gamma(G, F)}.$$

The *worst-case relative inefficiency* of the market game ϕ_m over the admissible environments E is

$$e^{wor}(\phi_m, E) = \sup_{(F, G) \in E} e(\phi_m, G, F).$$

It is the relative inefficiency of ϕ_m in that environment $(F, G) \in E$ that maximizes that inefficiency.

Given a set of admissible environments E and a set of market mechanisms M that are defined on E , a market mechanism $\Phi = \{\phi_2, \phi_3, \dots\} \in M$ is *worst-case asymptotic optimal* over E among market mechanisms in M if, for any $\Phi^* = \{\phi_2^*, \phi_3^*, \dots\} \in M$, a constant $\eta > 0$ exists such that

$$e^{wor}(\phi_m, E) \leq \eta e^{wor}(\phi_m^*, E)$$

for all $m \geq 2$. In other words, the market mechanism Φ is worst-case asymptotic optimal if no other market mechanism Φ^* exists such that, for each market size $m \in \{2, 3, 4, \dots\}$, the ratio between (i) its relative inefficiency in the environment within E it finds most difficult and (ii) Φ^* 's relative inefficiency in the environment that Φ^* finds most difficult is bounded by a finite constant. Said yet another way, Φ is worst case asymptotic optimal if no other mechanism has worst-case relative inefficiency that converges to zero at a faster rate than the double auction's worst case relative inefficiency.

Given this setup, the following theorem is true.

Theorem 2 *Satterthwaite and Williams (1999)*. *Let E be a set of admissible environments (F, G) such that (i) the uniform environment (G_U, F_U) is included, (ii) the densities (f, g) for every element of E satisfy $\underline{q} \leq f, g \leq \bar{q}$ for positive \underline{q}, \bar{q} , and (iii), for every environment $(F, G) \in E$, an equilibrium with positive probability of trade exists. Let M be the class of all interim individually rational and ex ante budget balancing market mechanisms. The double auction is worst-case asymptotic optimal over the environments E among the class M of market mechanisms.*

The proof of this theorem involves using mechanism design theory and, in particular, a new form for the incentive constraints that Williams (1998) developed.

We show that, for the uniform environment, the relative inefficiency of the constrained efficient mechanism, $\Phi^{\text{CE}} = \{\phi_2^{\text{CE}}, \phi_3^{\text{CE}}, \dots\}$, approaches 0 no faster than the relative inefficiency of the double auction, i.e., a $\xi > 0$ exists such that

$$e(\phi_m^{\text{CE}}, F_U, G_U) \geq \frac{\xi}{m^2}.$$

No other mechanism in the class of interim individually rational and ex ante budget balancing market mechanisms can converge faster than this on the uniform environment. Equilibria of the double auction converge at least that fast on the uniform environment and on any other environment that meets the conditions of Theorem 1.

While market mechanisms may exist that have a faster than quadratic rate of convergence on some environments within E , for such mechanisms the uniform environment is the “worst” case that is needed to make the theorem true. That the uniform environment plays the role of worst case is nice because the uniform case is anything but pathological. Our conjecture is that there are no mechanisms that converge at a faster than quadratic rate in any but pathological environments.

5 Conclusions

In this paper I have tried to make two main points. The first is that the double auction has remarkable properties when applied to a simple exchange market: in equilibrium the maximal misrepresentation of preferences is inversely proportional to the market’s size, the relative inefficiency is inversely proportional to the square of the market’s size, and no other mechanism dominates this rate of convergence to efficiency. These results provide a justification for price theory’s disregard of strategic behavior. The second, meta-point is that strategy-proofness in its exact form is too strong. There have been a few positive results such as Barberà, Sonnenschein, and Zhou’s (1991) “Voting by Committees” in which they found interesting and useful strategy-proof voting rules for a common group choice problem. Nevertheless the empirical prediction based on the many papers characterizing strategy-proof mechanisms must be that for almost all settings that remain to be explored no interesting strategy-proof mechanisms exist. Rather than continuing this program of characterizing strategy-proof mechanisms in more settings, it may be better to seek out approximately strategy-proof mechanisms.

The approach that I described above is only one possible direction for pursuing approximate strategy-proofness. For example, Kalai and Ledyard (1998) have shown how placing the group choice problem in a long run, repeated setting can be used to generate an approximately, or even exactly, strategy-proof mechanism. The work of Saari (1995) and Merlin and Saari (1997) on the geometry of voting suggests that it may be possible to determine the rate at which possibilities for manipulation decrease as the number of voters and alternatives increase. Bartholdi and Olin (1991) and Bartholdi, Tovey, and Trick (1992) use

complexity theory to show that computing whether manipulation is possible in a particular situation can be a NP-hard problem, thus suggesting that a convincing theory of approximate strategy-proofness might be based on bounded rationality.

References

- [1] Barberà, S. and M. Jackson. 1995. Strategy-proof exchange. *Econometrica* 63 (January): 51-88.
- [2] Barberà, S., H. Sonnenschein, and L. Zhou. 1991. Voting by committees. *Econometrica* 59 (May): 595-610.
- [3] Bartholdi, J. and J. Olin. 1991. Single transferable vote resists strategic voting. *Social Choice and Welfare* 8: 341-54.
- [4] Bartholdi, J., C. Tovey, and M. Trick. 1992. How hard is it to control an election? *Math. Comput. Modeling* 16: 27-40.
- [5] Blin, J. and M. Satterthwaite. 1975. On preferences, beliefs, and manipulation within voting situations. *Econometrica* 45 (May): 881-88.
- [6] Frank, R. 1991. *Microeconomics and Behavior*. New York: McGraw-Hill.
- [7] Gibbard, A. 1973. Manipulation of voting schemes: A general result. *Econometrica* 41: 587-602.
- [8] Gibbard, A. 1978. Straightforwardness of game forms with lotteries as outcomes. *Econometrica* 46: 595-614.
- [9] Gul, F. and A. Postlewaite. 1992. Asymptotic efficiency in large exchange economies with asymmetric information. *Econometrica* 60: 1273-92.
- [10] Hagerty, K. and W. Rogerson. 1987. Robust trading mechanisms *J. of Econ. Theory* 42: 94-107.
- [11] Hurwicz, L. 1972. On informationally decentralized systems. In *Decision and Organization*, ed. by B. McGuire and R. Radner, pp. 297-336. Amsterdam: North-Holland.
- [12] Hurwicz, L. and M. Walker. On the generic nonoptimality of dominant-strategy allocation mechanisms: A general theorem that include pure exchange economies. *Econometrica* 58 (May): 683-704.
- [13] Jackson, M. and A. Manelli. 1997. Approximately competitive equilibria in large finite economies. *J. of Econ Theory* 77: 354-76.
- [14] Kalai, E. and J. Ledyard. 1998. Repeated implementation. CMSEMS WP 1205R, Northwestern University, March. Forthcoming in *J. of Econ. Theory*.

- [15] McAfee, P. 1992. A dominant strategy double auction. *J. of Economic Theory* 56: 434-50.
- [16] Merlin, V. and D. Saari. 1997. Copeland method II: Manipulation, monotonicity, and paradoxes. *J. of Econ. Theory* 72 (January): 148-72.
- [17] Negishi, T. 1987. Tâtonnement and recontracting. In *The New Palgrave*, edited by J. Eatwell, M. Milgate, and P. Newman, V. 4, 589-596. New York: The Stockton Press.
- [18] Saari, D. 1995. *Basic geometry of voting*. Berlin: Springer-Verlag.
- [19] Satterthwaite, M. 1975. Strategy-proofness and Arrow's Conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *J. of Econ Theory* 10: 187-217.
- [20] Satterthwaite, M. and H. Sonnenschein. 1981. Strategy-proof allocation mechanisms at differentiable points. *Rev. of Econ. Studies* 48: 587-97.
- [21] Satterthwaite, M. and S. Williams. 1999. The optimality of a simple market mechanism. Xeroxed, February.
- [22] Schummer, J. 1997. Strategy-proofness vs. efficiency on restricted domains of exchange economies. *Social Choice and Welfare* 14: 47-56.
- [23] Williams, S. 1998. A characterization of efficient, Bayesian incentive compatible mechanisms. Xeroxed. Forthcoming in *Economic Theory*.
- [24] Wilson, R. 1987. Game-theoretic analyses of trading processes. In *Advances in Economic Theory, Fifth World Congress*, ed. by T. Bewley, pp. 33-70. Cambridge: Cambridge University Press.

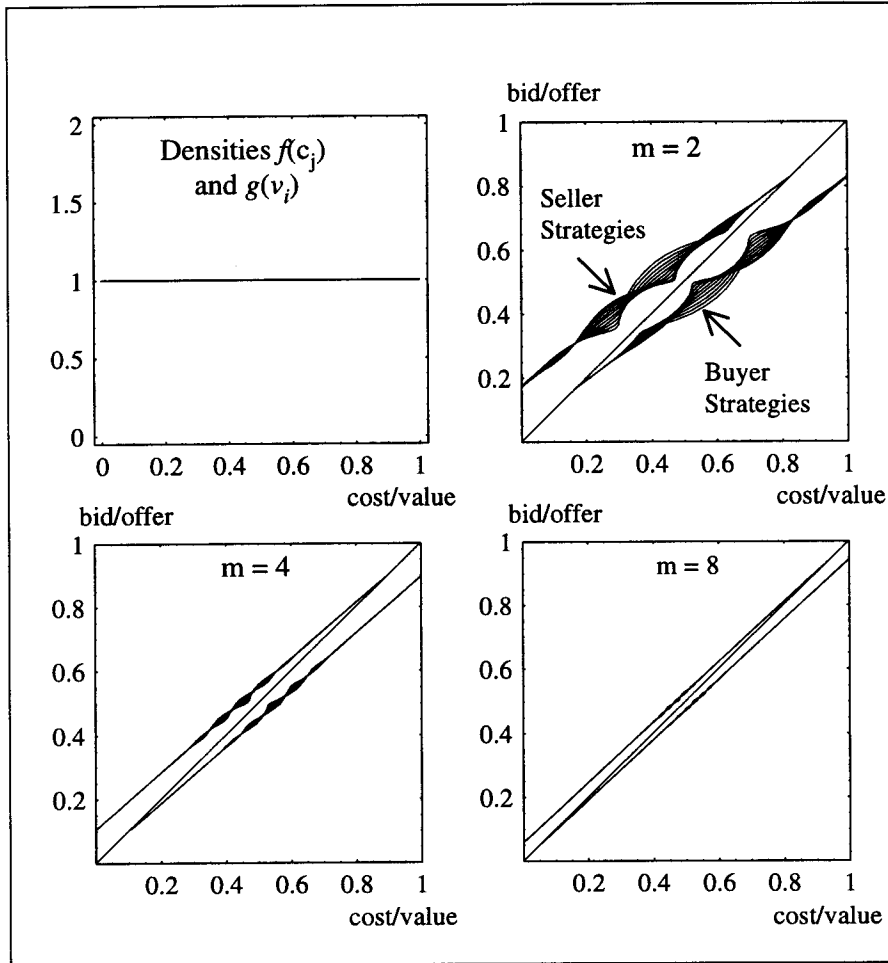


Figure 2: Bundles of equilibrium strategies for different market sizes m if values v_i and costs c_j are distributed uniformly on the unit interval.




Figure 3: Harberger triangle demonstrating that the relative inefficiency of the double auction converges to 0 quadratically.