Bank of Canada Banque du Canada

Working Paper 2002-33 / Document de travail 2002-33

Alternative Trading Systems: Does One Shoe Fit All?

by

Nicolas Audet, Toni Gravelle, and Jing Yang

ISSN 1192-5434

Printed in Canada on recycled paper

Bank of Canada Working Paper 2002-33

November 2002

Alternative Trading Systems: Does One Shoe Fit All?

by

Nicolas Audet,[1] Toni Gravelle,[2] and Jing Yang[3]

[1]University of Michigan
Ann Arbor, Michigan, United States

[2]International Monetary Fund
Washington, DC
(Formerly of the Financial Markets Department, Bank of Canada)
tgravelle@imf.org

[3]Financial Markets Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
jyang@bankofcanada.ca

Contents

Acknowledgements

We would like to thank Blake LeBaron and Richard Lyons for their comments. We thank Walter Engert for his continuous support and editorial assistance. We also would like to thank conference participants at the Bank of Canada and the 2002 NFA meeting.

Abstract

This paper examines the factors that lead liquidity-motivated investors to choose the type of market structure they prefer. We assume that investors can choose between a dealership and a limit-order-book market. This study builds a theoretical model for both the dealership and order-book markets and develops a numerical method to solve the Nash equilibrium strategies of heterogeneous market participants. We find that a dealership market would be preferred by investors in an environment where customer trading is relatively thin and correlated, and by investors who are subject to relatively large liquidity shocks.

JEL classification: G10, G14, G18
Bank classification: Financial markets

Résumé

Les auteurs examinent les facteurs qui amènent les investisseurs dont les transactions sont motivées par des chocs de portefeuille à privilégier une structure de marché particulière. Ils font l'hypothèse que les investisseurs ont le choix de recourir à un marché de contrepartie, dirigé par les prix, ou à un marché reposant sur la confrontation d'ordres à cours limité dans un carnet. Ils élaborent un modèle théorique pour ces deux types de marché, ainsi qu'une méthode de résolution numérique en vue de déterminer la stratégie d'équilibre de Nash d'opérateurs hétérogènes. Les auteurs constatent que les investisseurs préfèrent la structure du marché de contrepartie lorsque les transactions sont peu nombreuses et corrélées entre elles ou qu'ils subissent des chocs de liquidité relativement importants.

Classification JEL : G10, G14, G18
Classification de la Banque : Marchés financiers

1 Introduction

Financial markets are undergoing increasingly rapid and profound structural changes. The recent evolution of financial markets has been driven both by the unprecedented rapid creation and adoption of new information technologies and by changes in the regulatory framework.

Advances in information technology have led to the emergence of new electronic trading systems in various financial markets. For example, stock exchanges have faced increased competition from various new electronic trading systems, also known as alternative trading systems (ATSs) or electronic communication networks (ECNs). This is particularly true in U.S. equity markets, where nearly 40 per cent of the dollar volume of trading in Nasdaq shares has migrated to ATSs (Securities and Exchange Commission 2000).

As a result, ATSs represent an increasingly important competitive challenge to established trading venues, such as the New York Stock Exchange (NYSE) and Nasdaq. In explaining the success or failure of various trading venues to compete with each other, it is useful to examine how these trading venues differ in terms of their trading or market structure. Instinet and Island, two equity ATSs that offer an order-book trading mechanism, account for the largest proportion of trading of Nasdaq shares, with approximately 11 per cent and 14 per cent, respectively. The fact that these systems offer an order-book trading environment, whereas the Nasdaq offers a dealership system, has been cited as an important factor driving their success (Financial Services Authority 2000). Further supporting this hypothesis that a trading system's architecture affects a system's ability to compete for market share in the financial trading services industry is the fact that ATSs account for only 6 per cent of NYSE-listed share trading, where the NYSE is already set up as an order-book trading system supplemented by a specialist (monopoly market-maker). On the other hand, Conrad, Johnson, and Wahal (2002) show that over 90 per cent of the trading on crossing-type ATSs is for NYSE securities, which indicates that these systems may be able to succeed in attracting order flow for NYSE stock trading, where order-book ATSs have not.[1]

ATSs have failed to garner a significant proportion of trading activity away from traditional trading venues in other asset markets as well. Specifically, ATSs that offer an order-book trading mechanism have failed to become a significant force in the public segment of the foreign exchange or bond markets, where the vast majority of the public trading activity continues to take place in a dealership, quote-driven trading environment. On the other hand, those ATSs that represent the automation of the existing phone-based trading structure have tended to flourish in these markets. It may be that the success of these systems in terms of attracting trading volume is due to their ability to offer operational efficiency gains (lower trading costs) without having to significantly alter the market participants' desired trading structure. For example, the TradeWeb, a multi-dealer quote-driven bond trading system, has garnered an important amount of U.S. Treasury securities trading volumes. TradeWeb allows investors to solicit quotes electronically from a number of dealers and execute a trade with the dealer of their choice, which in essence replaces the need for investors to sequentially phone dealers for a quote on a bond and then trade with their preferred dealer (see Gravelle 2002). Similarly, order-book-style ATSs, such as Reuters Dealing

[1] In crossing systems, traders enter unpriced (buy and sell) orders, which are then crossed at prespecified times, at prices prevailing in the security's primary market. Conrad, Johnson, and Wahal (2002) argue that the higher the volume in the primary market and, in turn, the better the price discovery there, the more likely the crossing system is to attract order flow. Given their similar market-clearing structure, order-book ATSs do not offer much of a competitive advantage over the NYSE, in terms of their ability to provide price discovery. Instead, crossing systems compete with the NYSE in terms of their ability to more cheaply (measured in terms of lower price-impact costs) match larger uninformed buy and sell orders at the prevailing NYSE price.

2000-2 and EBS in the foreign exchange (FX) market and EuroMTS in the European bond market, have garnered a significant proportion of the trading activity in the interdealer segments of these markets (Bank for International Settlements 2001). These electronic broker systems also represent the automation of services offered by traditional voice-based interdealer brokers.[2]

Although the advent of new electronic trading technologies has the potential to allow new market structures to appear, providing market participants with a broader choice of trading environments, regulation also has the potential to affect certain features of the existing trading environment. One aspect of market architecture upon which regulation can have the greatest impact is transparency. Transparency is defined as the amount of prices and trading information that is disseminated to the broader market or to the public more generally. Given that trading venues compete, in part on the basis of differences in transparency, transparency rules applied to some markets and not to others will affect the degree of competition between trading venues and will create incentives for market participants or trading venues to engage in regulatory arbitrage (Sirri 2000 and Madhavan, Porter, and Weaver 2001). Bloomfield and O'Hara (2000) show that transparent trading venues are at a competitive disadvantage relative to trading venues that are more opaque. Although the application of common transparency rules across all marketplaces might serve to level the playing field and reduce the scope for regulatory arbitrage, it might also increase the implicit trading costs faced by certain investors, reducing their welfare. For example, large

[2]As argued in Lyons (2001), however, the electronic FX brokers provide an enhanced level of trade transparency to the dealers and, as such, would represent a change in trading structure of the interdealer segment of the FX market, rather than simple automation of an existing trading mechanism. We would argue that it is not clear whether it is the operational efficiency gains due to the automation (without significant alteration) of an existing mechanism, or whether it is the change in the structure of the market (offered by the greater transparency), that has led to the success of these FX ATSs.

liquidity-motivated public investors tend to enjoy lower trading costs when their orders are sent to more opaque dealership trading environments, than when their orders are routed to generally more transparent order-book trading venues (de Jong, Nijman, and Roel 1995, Ganley et al. 1998).

There is, therefore, a growing realization among regulators and competing trading venues that the design of the trading mechanism is an important determinant of a market's ability to successfully compete for order flow. To understand how different trading venues succeed in capturing trading activity in a particular security, one has to understand which attributes of the trading venue are most important for trading activity. Moreover, the ability of different trading mechanisms to offer lower trading costs to an investor depends, in turn, critically on the trading needs or characteristics of the investors themselves. Given that investors seek best execution of their trades, and that best execution encompasses traded price, market impact, immediacy, timing, anonymity, and commissions, it is not surprising that investors choose different trading venues based on how well each venue fulfils different combinations of these elements.

Our paper seeks to answer the following question: If customers could choose the market structure to trade in ex ante, which market structure would they choose? That is, what determines a market participant's choice of trading in an order-book market versus trading in a dealership market? Moreover, we seek to analyze how the trading environment and customer characteristics affect their choice of trading mechanism.

Despite the fact that new trading systems have attracted significant trading activity, little theoretical or empirical work has examined the considerations that determine a participant's choice of market mechanism. That investors are observed to send orders to trade a stock to different trading systems (Conrad, Johnson, and Wahal 2002) has not been the focus of much research. Most of the existing research concentrates on modelling one trading mechanism and often simply examines mechanism efficiency in terms of price discovery.

While it is highly desirable to examine which market structure best serves individual investor needs in terms of execution quality, it is generally very difficult to study both dealership and limit-order-book structures in a unified theoretical framework using standard market-microstructure modelling techniques. Empirical research that carries out intermarket comparisons is very difficult as well, because actual market structures are more complicated than the models on which empirical tests are based. Further, existing empirical research has carried out intermarket comparisons based largely on a measure of the bid-ask spread, which is only one dimension of execution quality. These studies fail, in general, to examine other measures of execution quality that could help explain why various market structures coexist and why customers send their orders to different systems. Empirical research is also impeded by a lack of detailed data surrounding events where the market structure has undergone a regime change. Such data would allow researchers to more directly test theoretical hypotheses; see Madhavan (2000).

To study the optimal market structure preferred by market participants, we develop an agent-based computational finance model of a dealership and limit-order market. This approach starts where theoretical market-microstructure models leave off, in that it allows researchers to examine questions that are analytically intractable in a purely theoretical construct. The agent-based financial market framework we develop is structurally grounded by an analytical model that guides the equilibrium behaviour of the artificial or simulated traders. One goal of this paper is to illustrate the applicability of the agent-based, artificial financial markets approach in the study of market design issues.[3]

[3]This methodology should not be confused with the experimental market studies that use human participants in a laboratory setting, such as those by Bloomfield and O'Hara (1999, 2000), Flood et al. (1999), and Theissen (2000).

The benefit of using an agent-based method is that it allows us to investigate questions related to optimal market structures in a unified analytical framework. Another benefit is that it allows for greater flexibility relative to analytical models, both in terms of how we specify the trading environment and how we model the characteristics of the market participants themselves. For example, to derive analytic solutions, many market-microstructure models assume that agents are limited to using linear pricing strategies and/or are risk-neutral. The agent-based simulation approach allows us to relax these and other restrictions.

The analytical framework upon which the agent-based simulation is constructed is based on a market-microstructure modelling approach. First, an institutional feature of many equity markets is that trading is organized as a continuous limit-order-book system. We model the limit-order-book market as a double auction, where market participants submit bid and ask orders to the system. The trading system then clears the market and determines the price. Each risk-averse market participant is rational and seeks to maximize expected utility from trading. As such, market participants realize, to the extent that they are the marginal buyer or seller in the system, the impact that their orders have on price, and they act strategically.

Second, we model the dealership trading architecture as a two-stage trading process. This captures a key institutional feature of multiple-dealer markets, such as FX and fixed-income markets (see Gravelle 2002). In the first stage, dealers trade with rational, liquidity-motivated public investors. These dealers (market-makers) commit to trade for any quantity at their posted bid and ask prices. After observing their liquidity shock, each public investor chooses the size of their order to submit against a dealer's quote. The risk-averse dealers subsequently retrade in the interdealer market via interdealer-broker (IDB) systems to lay off the inventory risk they obtained in the first stage of trading.

In the first stage of trading, dealers compete for customer order flow on price, à la

Bertrand competition, while second-stage interdealer trading is modelled as a limit-order trading mechanism. As noted in Viswanathan and Wang (2002), IDB systems in many respects resemble order-book trading systems. Therefore, the interdealer trading system is specified in an identical manner as the stand-alone limit-order market structure described above. The only modification is that it is only dealers that trade in this market, and they are motivated by their desire to lay off their unwanted inventory positions obtained in the first stage of trading. Dealers realize the impact that their orders have on their share of the total surplus among dealers and act strategically to maximize it. This two-stage trading process implies that the dealership market is actually a combination of two market structures. In turn, the dealer's quotation strategy in the first stage, when facing public investors, is a function of the trading environment they face in the second stage of trading.

We compare the welfare of rational public investors who trade in a dealership market with those who trade in the limit-order market. Although policy-makers and market designers are generally interested in the standards of execution quality that different market structures provide, we assume that the multidimensional nature of execution quality can, in the end, be summarized by its impact on investor welfare. We consider a framework where utility-maximizing public investors (customers) endogenously supply orders to the market based on a liquidity shock they receive just prior to entering the market. In essence, customer trading is motivated by the desire to share liquidity risk among a greater set of participants. Given that customers are identical except for their realization of the liquidity shock, the degree of customer heterogeneity is defined over their liquidity shocks. Adjusting the distribution from which the liquidity shocks are drawn allows us to vary customer characteristics related to order size. Customers are also characterized by the degree of correlation in their liquidity shocks. When this correlation is high, customer orders tend to be on one side of the market.

By varying customer characteristics, we examine a range of market structure issues. As noted in Viswanathan and Wang (2002) and Gravelle (2002), markets that primarily involve institutional traders who tend to generate large order flow, such as fixed-income and FX markets, are organized as dealership markets. Markets that handle primarily small orders generated mainly by retail stock-trading investors, such as certain Nasdaq stock trading ATSs and the downstairs segment of the Toronto Stock Exchange (TSX) and the Paris Bourse, are structured as limit-order-book markets. Our study investigates how the optimal market structure for customers whose order size is sometimes large and varies considerably may differ from that of customers characterized by relatively small homogeneous order flow.

We also examine customer choice of market architecture under different trading environments. The trading environment is defined over the thickness of the market, the number of market-makers, the degree of market-maker heterogeneity, and the risk-aversion differential between market-makers and customers. Different market structures could be better suited to overcome various coordination or trading frictions in financial markets. For example, a feature of fixed-income, and to a somewhat lesser extent, FX markets, which differs markedly from equity markets, is their thickness: the number of buyers and sellers trading in the market at any point in time.[4] As such, the observation that these markets are relatively thin might explain why the public trading segments of fixed-income and FX markets are structured as dealership systems.

Our findings suggest that the trading environment has an important impact on the optimal market structure. The public investor's choice to trade in a particular market will depend on the thickness of the market measured in terms of the number

[4]The fixed-income and FX markets are largely a wholesale market consisting of a relatively small number of large institutional investors. On the other hand, equity markets consist of thousands of traders, a large proportion of which are small retail investors.

of customers active in the market within a short time span. The dealership market structure is preferred by customers when there are few customers. As the number of potential public investors increases, there comes a point where the number of investors exceeds a critical threshold, so that these investors prefer to trade in the order-book market structure. We find that, as the number of dealers decreases or as the risk appetite of the dealers decreases, the risk-bearing capacity of the dealership market decreases, making this market structure less attractive to liquidity-motivated public traders. As a result, the limit-order-book market is more likely to prevail. Increasing dealer heterogeneity is found to increase the likelihood that the dealership system will prevail. Customers who are subject to larger and more volatile liquidity shocks or whose orders are correlated with others will prefer to trade in a dealership system. This is consistent with the observed regularity in which capital markets, where trading is infrequent or dominated by a small number of large institutional investors, tend to be organized as dealership markets.

This paper is organized as follows. In section 2 we review some of the related market-microstructure literature and agent-based computational finance literature. The analytical specification of the trading environment is presented in section 3. As with most market-microstructure models, both the dealership and order-book market structures are specified as simultaneous or one-shot games, in which agents must choose their strategies (actions) to maximize their utility while taking the other agent's strategies into consideration. The second part of section 3 describes how the agent-based framework is used to solve for the equilibrium solution of the analytical model. Section 4 reports our computational findings, and section 5 reviews some of the limitations of our current framework. We also provide suggestions on how future extension of our framework would allow customers to have private information about the securities' fundamental value. Appendix A describes the numerical procedure used to find the optimal strategies for each dealer and the Nash equilibrium.

2 Related Literature

2.1 Market microstructure

Although the market-microstructure literature has grown tremendously since the early 1980s, relatively little theoretical or empirical work has been done to determine the factors that affect a market participant's choice of market mechanism.[5] Recently, a small but growing number of papers have attempted to study optimal market design issues by modelling a trading mechanism and then adjusting it, in turn creating a modified trading structure. For example, a number of authors have recognized the importance of market design and have sought to explain why markets with differing levels of transparency coexist (Lyons 1996, Flood et al. 1999, Bloomfield and O'Hara 1999, and Naik, Neuberger, and Viswanathan 1999).[6] In general, their papers show that a trade-off exists between market or execution quality (often measured in terms of market liquidity) and market efficiency (measured in terms of effective price discovery), and that this trade-off may explain why market structures with differing levels of price transparency coexist. Of particular relevance to our study is the theoretical research that explicitly examines trader welfare across different market structures, such as Vogler (1997), Naik, Neuberger, and Viswanathan (1999), Saporta (1997)

[5] O'Hara (1995), Madhavan (2000), and Lyons (2001) provide excellent overviews of the existing literature.

[6] Although there is a vast literature examining market transparency, most researchers examine how various levels of pre-trade transparency affect the price-discovery process and not why markets with different levels of transparency coexist. This literature on market transparency tends to ignore issues related to implicit trading costs or, more broadly, execution quality issues. Moreover, it is generally assumed that highly price-transparent trading mechanisms increase the market participants' ability to extract information from outstanding quotes, and ensure that prices reflect a maximum of available information. As Flood et al. (1999) found, however, greater pre-trade transparency may in fact restrict the price-discovery process.

and Viswanathan and Wang (2002).[7]

Vogler (1997) develops a model that compares dealership markets with auction markets. The key and realistic feature of Vogler's model is the fact that the price the dealer offers to the liquidity-motivated public investor (in the first stage of trading) will depend on the dealer's ability to manage their inventory position in the second stage, where they are assumed to trade via an auction mechanism. Vogler shows that if customer trading is motivated by liquidity shocks, a dealership market provides customers with a better price than an auction market. This result holds as long as the market-maker can rebalance their inventory in the interdealer market. Vogler models the first-stage public trading as a Bertrand competition game. By contrast, the auction market is modelled as a quantity competition game. As Vogler notes, the limitation of this paper is the Walrasian market-clearing mechanism assumed in both the interdealer trading or second stage of the multiple-dealer market and in the auction market. This mechanism assumes away execution risk that actually occurs in pure order-book markets. Moreover, in Vogler (1997), as in most of the literature, customers are exogenous liquidity traders. In our study, we endow each customer with a utility function that they seek to maximize. This makes their trading decisions endogenous, which is not the case in most of the literature.

As in Vogler (1997), Naik, Neuberger, and Viswanathan (1999) develop a two-stage theoretical model that resembles most multiple-dealer markets. In Naik, Neuberger, and Viswanathan (1999), second-stage trading assumes that dealers split their order among the remaining dealers who were not privy to the customer order, which contrasts with Vogler's model, where the dealer trades with only one of the remaining

[7]Empirical work that compares execution quality across different market structures includes de Jong, Nijman, and Roel (1995), Keim and Madhavan (1997), Bessembinder (1999), Smith et al. (2000), and Madhavan, Porter, and Weaver (2001). Bloomfield and O'Hara (2000) and Theissen (2000) examine market structure competition issues using an experimental markets methodology.

dealers. Moreover, Naik, Neuberger, and Viswanathan (1999) examine how customer welfare is affected under two trade-disclosure regimes—one where post-trade information is made public and one where it is not—whereas Vogler was interested in the efficiency of the two-stage trading process relative to a one-period standard auction. An innovation on Vogler's (1997) study, where public investors are assumed to be risk-neutral, is that Naik, Neuberger, and Viswanathan (1999) model the investor and market-maker as being risk-averse, and make endogenous the investor's trading decision, which is conditional on the dealer's price schedule. Naik, Neuberger, and Viswanathan show that, in some circumstances, dealership markets that require full and prompt disclosure of investor-to-dealer trade details may reduce the welfare of risk-averse investors. The likelihood that prompt post-trade disclosure impairs investor welfare depends in part on the size of the endowment or liquidity shock received by the investor. The authors show that an increase in the level of post-trade transparency works against the execution of trades for investors hit by large liquidity shocks, but tends to benefit those subject to small liquidity shocks.

Saporta (1997) models a risk-averse dealer's choice between trading on two different interdealer trading mechanisms. She develops a three-stage model, that is similar in structure to that of Naik, Neuberger, and Viswanathan, to study the dealers' choice of trading bilaterally with other dealers or placing an order through an IDB, and assesses the implication of a multiple interdealer trading mechanism for the liquidity-motivated customers. In modelling the direct or bilateral interdealer trading mechanism, Saporta assumes that the dealers compete, in Bertrand fashion in terms of their price schedules. The IDB trading mechanism is modelled as a continuous order-driven market where dealers take into account their orders' impact on the market-clearing price. Saporta allows for the endogenous entry of dealers into the market-making of the risky security, which contrasts with the related multiple-dealer literature described in this section. She shows that investors subject to large liquidity needs are made better off when the dealers have the choice of using two different interdealer trading structures. When dealers become more risk-averse or when investor liquidity trading increases, Saporta's model predicts that interdealer trading is likely to be brokered rather than direct. Her findings also indicate that an increase in transparency causes the interdealer trading environment to become less efficient in terms of its inventory risk-sharing capabilities, thus increasing the number of market-makers required to make interdealer markets viable.

Recent theoretical work by Viswanathan and Wang (2002) is closest to our paper, in that it seeks to answer the following question: faced with a choice between different market structures, which one would a risk-averse customer choose? Specifically, Viswanathan and Wang examine how customers fare across three market structures: limit-order-book, dealership, and hybrid market structures, the latter being a combination of both the order-book and dealership structures, in which large orders are constrained to be sent to the dealership structure. The authors model these three different structures using three different market-clearing mechanisms: discriminatory price auction, uniform price auction, and a hybrid of those two auction mechanisms. They relate the discriminatory and uniform market-clearing structures to order-book and dealership market structures, respectively. Viswanathan and Wang show that risk-neutral customers will always prefer to participate in a discriminatory price auction, the order-book structure. A dealership structure is preferred by risk-averse customers when the number of market-makers is sufficiently large. They also find that for risk-averse customers, the hybrid market structure can dominate the dealership and order-book structures. In contrast with the papers discussed above, Viswanathan and Wang do not model second-stage trading in the dealership market.

The theoretical market microstructure aspects of our study differ from the previous work in several important respects. First, our paper differs from all papers but the Naik, Neuberger, and Viswanathan paper, in that we allow customers to rationally choose the size of the order they submit to the market. Second, we use a double-auction market-clearing framework in modelling the order-book structure, which is a more realistic characterization of this type of market. This framework implies that customers trade with each other and can be both suppliers and demanders of liquidity. We also innovate on the extant double-auction literature by relaxing the constraint that limit-order size be exogenously fixed.[8] The order-book market in Viswanathan and Wang (2002) is a one-sided auction framework where customers trade with dealers who supply a bid/offer schedule. (Note that, in Vogler, customers trade in a Walrasian auction setting, and in Saporta (1997) and Naik, Neuberger, and Viswanathan (1999), the customers cannot choose to trade in an order-book setting.) Third, in modelling the dealership market, we model interdealer trading as a double-auction rather than a Walrasian or simple order-splitting structure as in Vogler (1997) and Naik, Neuberger, and Viswanathan (1999), respectively. Fourth, we allow traders' strategy functions to be non-linear, which are based on a flexible-form artificial neural network (ANN). Linear response functions dominate the market-microstructure literature.

2.2 Agent-based computational finance

Agent-based computational finance is a branch of economics that models financial markets from the "ground up," in which the dynamic interaction of heterogeneous agents is examined. In these models, agents maximize their profits or wealth based on trading strategies (see LeBaron 2000, 2001, 2002 for reviews of this literature). By starting with computer simulations of strategies of many individual adaptive agents who trade in a given institutional/market structure, and then analyzing the resultant macroeconomic outcomes, this approach can be used to analyze the aggregate dynamics of various economic systems, including the price dynamics typically observed

[8] See Werner (1997), and Rustichini, Satterthwaite, and Williams (1994) for examples of fixed-order size double-auction models.

in financial markets.

The agent-based approach is a subset of a larger and increasingly important area of economics often called computationally based non-linear dynamics (Dechert and Hommes 2000 and Brock 2000), which has its roots in the analysis of economic fluctuations (i.e., business cycles), and has been applied to examine issues related to the evolution of economical systems and social norms, learning, auction markets, sustainable natural resource use, labour markets, and network economics, to name a few.[9]

Most of the agent-based research on artificial markets has been developed in the hope of explaining empirical puzzles that are not consistent with the efficient-markets hypothesis, one of the pillars of financial economic theory. That is, empirical evidence indicates that financial markets are subject to volatility clustering, overshooting/undershooting (relative to fundamentals), market crashes or crises (i.e., fat tails), and speculative bubbles. Additionally, asset prices have been found to be predictable in the sense that past price information helps predict future asset prices (i.e, simple technical analysis is profitable).

The efficient-markets hypothesis predicts that past information cannot help predict future prices, and markets are assumed to be free of (endogenous) dynamics that are not reflective of exogenous changes in fundamentals. The efficient-markets hypothesis stems from theoretical equilibrium models based on the assumed presence of homogeneous rational traders that possess all relevant information. More generally, standard general-equilibrium models assume the presence of a representative agent, which in a fundamental way precludes the possibility of trading, as there is in essence only one type of agent. Equilibrium prices adjust instantaneously to new equilibrium levels without trading after the release of fundamental information in these models.

[9] Leigh Tesfatsion maintains a Web site that archives the various applications of agent-based computational methods, at http://wwww.econ.iastate.edu/tesfatsi/ace.htm.

Thus another empirical puzzle resulting from these models is to understand why so much trading is taking place in financial markets. As such, modelling the trading activities of multiple heterogeneous agents who are rational or boundedly rational has been suggested as a way of developing an understanding of the complex price dynamics and the associated trading volume that the equilibrium efficient-markets models cannot explain.

Although the use of heterogeneous agents in theoretical market-microstructure research, where three types agents are typically defined (market-maker, and informed and uninformed investors), is not new, the degree of heterogeneity in these models is in general quite limited relative to the agent-based approach. The agent-based framework allows us to examine a much richer extent of heterogeneity than is analytically tractible. Moreover, in agent-based computational models, the individual trading strategies are themselves often non-linear, as in our study, and are allowed to change over time.

Given that designers of the agent-based artificial markets have primarily been interested in explaining the puzzling out-of-equilibrium price dynamics, they have emphasized learning behaviour and the evolution of the trading strategies of agents in their artificial financial markets. Traders are assumed to use strategies (typically, forecasting rules) based on the learned relation between prices and market information, and their individual strategies or behaviours are allowed to evolve in response to the past performance of those strategies. Moreover, as a consequence of their focus on the efficiency properties of asset-price dynamics, designers of agent-based models have assumed simple trading structures, such as Walrasian auction. As such, how the traders interact (or actually trade) in agent-based financial market models often does not closely match how real markets actually operate. One exception is Yang (2002), who introduces a double-auction trading system and narrows the gap between agent-based artificial markets and real financial markets. Hence, she makes it possible to compare the simulation results with real tick-by-tick data.

Given our focus on comparing the welfare properties of two trading systems, it is important that the market outcomes that we analyze result from the optimal strategies of agents. Therefore, in contrast with the existing agent-based literature, using a Nash equilibrium framework we are able to examine the welfare implications of changes in market structure based on the rational behaviour of agents. This approach de-emphasizes asset-price dynamics while designing market-clearing mechanisms that more closely match those observed in real financial markets. In so doing, we develop an agent-based framework that is more closely aligned with the theoretical models in the market-microstructure literature discussed in section 2.1, except that we allow for greater heterogeneity, and use non-linear trading strategies and a complex market-clearing mechanism.

Nicolaisen, Petrov, and Tesfatsion (2001) also focus on trading-mechanism design issues. Using agent-based simulation, they study the efficiency of the discriminatory double-auction trading mechanism in a computational electricity market. Their results suggest that the discriminatory double auction reliably delivers high market efficiency when traders are rational. This result may not be robust, however, when the learning behaviour is not well considered; when, for example, traders do not respond to zero-profit trading outcomes. As discussed in more detail in section 3, the way we set up our payoff function differs from the method used by Nicolaisen, Petrov, and Tesfatsion. They use a "reinforcement learning algorithm" to overcome the typical payoff-function discontinuity problem observed (but rarely discussed) in most agent-based financial market settings, while we apply a novel smoothing-parameter approach that modifies the payoff function directly.

3 Model Description

The model we consider assumes the existence of two types of agents: customers and dealers. Both possess a portfolio composed of a risk-free and a risky asset; the liquidation or fundamental value of the risky asset is normally distributed and the distribution is common knowledge to all participants.

Moreover, all agents manage their portfolio to optimize a negative exponential utility criteria. Thus, all agents in our model will behave as rational traders. In addition to the fact that dealers are obliged to make the market and thus to provide liquidity for customers, the two agent types differ only in terms of the market mechanisms they can access to rebalance their portfolio. This feature of our model will allow us to change more freely the configuration of the markets without altering the fundamental characteristics of those agents. Consequently, the framework's flexibility allows us to use comparative statistics to compare different types of market organization in terms of the welfare of the participants and the liquidity of the market.

In section 3.1 we describe the models for two different markets: a limit-order-book market and a dealership market.

3.1 Limit-order market model

3.1.1 Trading environment

We assume that M customers (public investors) trade a risky asset in the market. They are identical in every way except for the realization of their exogenous inventory shock, $(X_i)_{i\in\mathcal{M}}$. The liquidation value of the risky asset is F. The random variables F and $(X_i)_{i\in\mathcal{M}}$ are independent normally distributed and the distribution is common knowledge to all participants.

Consider the following sequence of events. First, each customer receives nonpayoff-related information, inventory shock X_i. Each customer then simultaneously submits a bid and an ask price-quantity pair, $\left(B_i, Q_i^B\right)$ and (A_i, Q_i^A), respectively, conditional on their realization of the shock. A bid quote indicates that if the market-clearing price is below the bid, B_i, the customer agrees to buy Q_i^B units of the risky asset at the market price. Also, the ask quote indicates that if the market-clearing price is above the ask, A_i, the customer agrees to sell at the market price up to Q_i^A units of the security. Each customer knows the distribution from which liquidity shocks are drawn, but they do not observe the other customers' shock.

The market-clearing price is determined by a double-auction system. The indicated trades are executed at this market-clearing price. Since the market-clearing price is determined after the customers have submitted their orders, each customer will realize that their bid-ask orders may influence the clearing price to the extent that they are the marginal buyer or seller in the auction. Finally, the value of the risky asset is revealed to the market participants. Their wealth can therefore be evaluated as well as their utility.

The problem faced by each customer is to find a strategy that maximizes their expected utility. We assume that all market participants have exponential utility functions defined over wealth W with risk-aversion coefficient λ. Thus the utility function of customers is given by

$$U_i = -\exp\left(-\lambda_i W_i\right), \tag{1}$$

and the wealth of a customer can be written as

$$W_i = X_i F + Q_i \left(F - P\right), \tag{2}$$

where Q_i is the executed quantity.

Table 1: Order Matching and Trading Outcome

Participants	Inventory shock	Bid order (P, Q)	Ask order (P, Q)	Quantity executed	Wealth (F=1)
Player 1	-0.5	($0.7, 2)	($0.8, -1)	2	0.5
Player 2	0.6	($0.4, 2)	($0.5, -3)	-1(partial)	0.1
Player 3	0.4	($0.1,0.5)	($0.2, -1)	-1	-0.1

3.1.2 Market-clearing system: double auction

We model the limit-order-book market as a k-double auction (k-DA). The standard set-up of k-DA is fairly restrictive, in the sense that the order size is exogenously fixed (Werner 1997, Rustichini, Satterthwaite, and Williams 1994). In this study, we generalize this k-DA set-up to allow each order size to be endogenously decided by market participants. This market-clearing mechanism has some intuitive appeal, as the quantities demanded for prices above this clearing price are sequentially matched with quantities supplied at prices below it until there are no orders at which the purchase price exceeds the selling price. In other words, orders are matched sequentially until the aggregate demand curve lies below the supply curve. This auction-clearing system implies that all trades clear at the same price. We start with a numerical example to illustrate this double-auction clearing system, and then provide a formal definition.

Table 1 gives an example of a trading outcome for a market that consists of three participants. A negative quantity indicates a sell order.

Each participant starts with a zero endowment of cash and a certain amount of a risky asset. The size of the risky asset is given by the liquidity shock. The initial liquidity shock for players 1, 2, and 3 is given in the second column. Conditional on their initial shock, each customer chooses an action according to their strategy function. An action consists of a bid price-quantity pair and an ask price-quantity pair. For example, after receiving a negative shock, participant 1 tends to rebalance their portfolio by posting a greater buy order. Actions that each participant takes are shown in columns 3 and 4. Then the market-clearing price in a k-DA mechanism is set to 0.5. All the bid prices that are higher than (or equal to) 0.5 and all the ask prices that are below (or equal to) 0.5 will be executed. In this case, the allocation is carried out as far as possible by assigning priority to sellers whose offers are smallest and buyers whose bids are largest. Partial execution may exist at the market-clearing price. In this case, participant 2 had a partial execution of their ask order. Given the initial allocation of the risk-free and risky assets, the transaction, and the realization of F at the end of the game, the wealth of each participant can be evaluated as shown in the last column of Table 1.

A market-clearing price in a k-DA can be formally defined as follows. In the standard k-DA set-up, the quantity associated with a quote is exogenously fixed at a value Q.[10] Therefore, a quote can be completely described by (A_i, B_i). The equilibrium price, P, is determined by the following method: let $\mathcal{P}$ be the set of all prices that clear the excess demand function and the equilibrium price be a convex combination of those two prices,

$$P = k \sup \mathcal{P} + (1-k) \inf \mathcal{P}, \tag{3}$$

where each choice of $k \in [0,1]$ defines a different mechanism.

In our set-up, we generalize this fixed-order size k-DA to variable quantities; as a result, $\mathcal{P}$ defined above might lead to an empty set, since we are looking for a price P such that excess demand is positive for all the prices that are less than P and negative

[10] Moreover, in the standard set-up, the auction is one-sided, in the sense that all participants are either buyers or sellers, but not both.

for all the prices that are greater than P. Moreover, we want to have a definition that yields a unique solution even if there are multiple prices that solve the excess demand. Therefore, we redefine

$$XD(p) = \sum\nolimits_{i=1}^{M} Q_i^B \, I_{(p \leq B_i)} - \sum\nolimits_{i=1}^{M} Q_i^A \, I_{(A_i \leq p)}, \tag{4}$$

where $I_{(p \leq B_i)}$ and $I_{(A_i \leq p)}$ are the indicator function, which is equal to one if the condition is true, and zero otherwise. The boundaries P^+ and P^- will be defined by

$$\begin{aligned} P^+ &= \sup\{p : XD(p) > 0\}, \\ P^- &= \inf\{p : XD(p) < 0\}. \end{aligned} \tag{5}$$

This price is defined by

$$P = k\,P^+ + (1-k)\ P^-, \tag{6}$$

using the new definition.

3.1.3 Numerical method to search for a Nash equilibrium

The market participant's strategy function maps information to an action, such as a bid-offer schedule $\left(B_i, Q_i^B, A_i, Q_i^A\right)$. The information used by participant i is called a liquidity shock, X_i. Therefore, the strategy function, s_i, is a mapping from X_i to $\left(B_i, Q_i^B, A_i, Q_i^A\right)$. A *best response* of participant i is a strategy that maximizes participant i 's utility. Given the participant's initial liquidity shock and other participants' strategies, a set of strategies where every participant plays their best response is called a *Nash equilibrium* for the auction.

We propose a numerical method to solve for the Nash equilibrium. The method is described in detail in Appendix A. Here we discuss only two necessary conditions to apply this numerical method.

The first necessary condition to apply the numerical method is that the utility function is continuously differentiable. The utility function is continuous in wealth but not with respect to the parameters of the strategy function.

For example, in this double-auction game, a customer's order is not executed when his/her bid price lies below the market-clearing price. As a result, there is no change in the customer's utility function. As the customer's bid increases slightly, nothing happens to the utility function until the moment the bid is greater than the market-clearing price. At that moment, this customer becomes a winner of the auction, and the order is executed. Therefore, there is a discrete change in the customer's utility function.

This discontinuous feature of the payoff function prevents the use of usual optimization methods. We develop a smoothing method to modify the payoff function such that the payoff function is differentiable. When the smoothing parameter approaches a limit, the modified payoff function converges to the original discrete payoff function. In our simulation, the criteria of convergence is that the difference of the two functions is 10^{-5}. Using this smoothing method, each participant receives a proportion of the auction goods, and the proportion decreases as the difference between the participant's bid and the winning bid increases. This proportion is controlled by the smoothing parameter. The greater the parameter, the larger the proportion of the goods the winning bid gets. Therefore, the winning bidder gets the whole auction at the limit, since the payoff function converges to the original step payoff function.

The second condition needed for a numerical solution is the parameterization of a strategy function. Since it is very difficult to assume the functional form of the mapping between a liquidity shock and an action, we use the following flexible form ANN for an individual's strategy function:

$$ANN(X_i, \theta_i) = \theta_{i,0} + \sum\nolimits_{k=1}^{H} \theta_{i,k} \ \varphi\left(\theta_{i,(H+1)+k} + \theta_{i,(2H+1)} \ X_i\right). \tag{7}$$

Therefore, the input of this strategy function is liquidity shock X_i. The output of this ANN is a vector of four choice variables$\left(B_i, A_i, Q_i^B, Q_i^A\right)$, and a strategy function is characterized by the vector of parameters θ_i, which contains $\theta_{i,0}$, and $\theta_{i,k}$. H is the number of hidden nodes. The *activation function* φ is chosen as the tanh function:

$$\varphi\left(x\right) = \tanh\left(x\right) \tag{8}$$

We slightly modify the first two choice variables, B_i and A_i, to the midpoint of the bid and ask price, $\frac{B_i+A_i}{2}$, and the spread of the bid and ask, $A_i - B_i$. This modification is to guarantee that an ask price is always higher than a bid price. The resulting specification for each market participant's strategy function is:

$$\begin{bmatrix} \frac{B_i+A_i}{2} \\ A_i - B_i \\ Q_i^B \\ Q_i^A \end{bmatrix} = \begin{bmatrix} ANN\left(X_i, \theta_i^{mid}\right) \\ \phi\left(ANN\left(\theta_i^{spread}\right)\right) \\ \phi\left(ANN\left(X_i, \theta_i^{Q^B}\right)\right) \\ \phi\left(ANN\left(X_i, \theta_i^{Q^A}\right)\right) \end{bmatrix}. \tag{9}$$

Therefore, θ_i is $\left(\theta_i^{mid}, \theta_i^{spread}, \theta_i^{Q^B}, \theta_i^{Q^A}\right)'$. $\theta_i^{mid}, \theta_i^{spread}, \theta_i^{Q^B}$, and $\theta_i^{Q^A}$ are parameters in vector θ_i that are used to generate the midpoint of the bid and ask, spread, and bid depth and ask depth, respectively. Notice that ϕ is applied to three outputs of the ANN to guarantee that the bid-ask spread and the bid and ask depth are positive. The function ϕ is given by $\phi\left(x\right) = \log\left(1 + \exp\left(x\right)\right)$.

Figure 4 depicts the bid and ask for a typical customer as a function of their liquidity shock, X_i, in their risky-asset position. At the beginning of each period, customers receive different liquidity shocks, and therefore post different bid and ask prices. The realizations of market-clearing prices, which are determined using equation (4), are represented by the dots in the middle part of the graph. The upper and lower curves are ask and bid prices for given liquidity shocks for a representative customer. Since an order can be executed only if the ask (bid) price is greater (lower) than the market-clearing price, the figure shows that a customer tends to post a low ask price to try to sell part of their inventory for rebalancing their portfolio if hit by a positive liquidity shock.

See section 6.2 of Appendix A for more details on the smoothing method and parameterization of the strategy function.

3.2 Dealership market

3.2.1 Trading environment

This section describes a trading model with M customers indexed by j and N dealers indexed by i. They all agree on the distribution of the final value of the risky asset. Trading takes place in two stages. In the first stage, dealers trade with customers. Dealers post bid and ask quotes simultaneously prior to the arrival of customer orders. The bid and ask quotes are valid for any order size. A customer receives an inventory shock, X_j, prior to observing dealers' quotes. Upon observing those quotes, the customer chooses the quantity to trade with dealers by maximizing their expected utility. Customer order flow is price elastic, in the sense that it decreases with the bid-ask spread that dealers quote. The trade size between the customer and dealer is private information. Therefore, the customer order flow determines the distribution of inventories among dealers at the end of the first stage.

In the second stage, dealers trade strategically in the interdealer market to reallocate the inventories that they carried from the first stage. Dealers can submit a bid quote and an ask quote through the double-auction system. The market-clearing price in the second stage is determined in the same manner as in an order-book market. Orders are matched sequentially starting from the highest buy order to the lowest ask order. This matching process continues until there are no more orders for which the bid price is higher than or equal to the ask price. In the end, a fundamental value, F, is revealed and every dealer's wealth is evaluated according their trading outcomes and the fundamental value of the risky asset.

In the first stage, or public trading round, each customer does not observe quotations from *all* N competing dealers. The reason is as follows. After observing the quotes, the customer decides whether to split their order between dealers or trade the whole order as a block to one dealer who quotes the best price. If all dealers start with the same zero inventory position, allowing customers to split their order equally among all the dealers results in an optimal risk allocation, which means that no dealer has an incentive to trade in the interdealer market in the second stage. However, allowing for different initial dealer inventories would motivate interdealer trading even if the customer equally split their order. We explore this variation in one of our experiments. For most of our experiments, we allow customers to observe quotes from *different subsets* of N competing dealers and a customer trades only with the dealer who offers the best observed price without splitting their order. In fact, it is a defining characteristic of the dealer market that dealers compete for the whole order from a customer.

3.2.2 Trading strategies at stage 1 and stage 2

Customer strategy in stage 1 Each customer, j, observes a randomly chosen set of k dealers, Δ_i, and decides on the quantity, C_{ij}, that they wish to trade with dealer i. The total amount traded by customer j is $C_j = \sum_{i \in \mathcal{N}} C_{ij}$. The dealers then trade among each other by participating in a double auction, where they submit a bid quote

$\left(B_i^2, Q_i^{B2}\right)$ and an ask $\left(A_i^2, Q_i^{A2}\right)$. A quantity, Q_i, will be executed at a price, P, and, finally, a fundamental value, F, will be revealed.

In the first stage, each customer, j, observes a set of quotes, $\left(B_i^1, A_i^1\right)_{i\in\Delta_j}$, randomly drawn from k dealers, and the set of dealers is denoted Δ_i. Note that the maximum bid of that subset is $\overline{B}_j^1$ and the minimum ask is $\underline{A}_j^1$, and the corresponding dealers are δ_j^A and δ_j^B. Each customer, j, decides the quantity, C_j, that they wish to trade by maximizing the expected utility of wealth over C_j, where utility is given by

$$MaxE(U_j^C) = -\exp\left(-\lambda_j^C\, W_j^C\right),$$

and wealth is given by

$$W_j^C = \begin{cases} X_j\, F + C_j\, \left(F - \underline{A}_j^1\right) & if \quad C_j > 0 \\ X_j\, F & if \quad C_j = 0 \\ X_j\, F + C_j\, \left(F - \overline{B}_j^1\right) & if \quad C_j < 0 \end{cases}. \tag{10}$$

X_j is known at the moment C_{ij} is decided and F follows a normal distribution with mean 0 and variance σ_F^2. Therefore, the optimal value of C_{ij} can be solved as

$$C_j = \begin{cases} \dfrac{-\overline{B}_j^1}{\lambda_j^C\, \sigma_F^2} - X_j & if \qquad P_j^R < \overline{B}_j^1 \\ 0 & if \quad \overline{B}_j^1 < P_j^R < \underline{A}_j^1 \\ \dfrac{-\underline{A}_j^1}{\lambda_j^C\, \sigma_F^2} - X_j & if \qquad \underline{A}_j^1 < P_j^R \end{cases}, \tag{11}$$

where the reservation price, P_j^R, of customer j is given by

$$P_j^R = -\lambda_j^C\, \sigma_F^2\, X_j. \tag{12}$$

Therefore, the optimal value of C_{ij} is

$$C_{ij} = \begin{cases} C_j & if \quad i = \delta_j^B \text{ and } C_j < 0 \\ C_j & if \quad i = \delta_j^A \text{ and } C_j > 0 \\ 0 & otherwise. \end{cases} \tag{13}$$

See Appendix A for more details.

Dealer's strategy in stage 1 and stage 2 Dealers provide one bid-ask quote $\left(B_i^1, A_i^1\right)$ upon a customer's request and then receive the customer's order flow at the prices they quote. In the second stage, dealers quote a bid and ask price-quantity pair, $\left(B_i^2, A_i^2, Q_i^B, Q_i^A\right)$, and trade in an interdealer market to share their inventory risk among each other. The dealer's objective is to maximize their utility over these six choice variables. Using notation that is similar to what we used for the order-book market, we can write the objective of the dealers as

$$MaxU_i^D = -\exp\left(-\lambda_i^D W_i^D\right), \tag{14}$$

subject to

$$W_i^D = \sum\nolimits_{j \in \mathcal{M}} C_{ij} \left(I_{(C_{ij}>0)}\, A_i^1 + I_{(C_{ij}<0)}\, B_i^1\right) - Q_i\, P + F\left(Q_i - \sum\nolimits_{j \in \mathcal{M}} C_{ij}\right). \tag{15}$$

For the strategy of stage 1, $\left(B_i^1, A_i^1\right)$, a single parameter is used for the bid-ask spread, $A_i^1 - B_i^1$. The midpoint of the bid-ask spread, $\frac{B_i^1 + A_i^1}{2}$, is a function of the initial inventory shock of dealers, which is set at zero in the baseline experiment we examine. This constraint guarantees that, for any customer, we will have $\underline{A}_j^1 > \overline{B}_j^1$, and it therefore guarantees that no arbitrage opportunities are available to a customer. In stage 2, since the interdealer market is modelled as a double-auction limit-order-book market, a dealer's trading strategy in this stage is identical to the one we described in the order-book market. A dealer's strategy function is summarized in the following equation,

$$\begin{bmatrix} \frac{B_i^1+A_i^1}{2} \\ A_i^1 - B_i^1 \\ \frac{B_i^2+A_i^2}{2} \\ A_i^2 - B_i^2 \\ Q_i^B \\ Q_i^A \end{bmatrix} = \begin{bmatrix} ANN\left(X_i^D, \theta_i^{mid1}\right) \\ \phi\left(ANN\left(\theta_i^{spread1}\right)\right) \\ ANN\left(C_i, \theta_i^{mid2}\right) \\ \phi\left(ANN\left(\theta_i^{spread2}\right)\right) \\ \phi\left(ANN\left(C_i, \theta_i^{Q^B}\right)\right) \\ \phi\left(ANN\left(C_i, \theta_i^{Q^A}\right)\right) \end{bmatrix}. \tag{16}$$

Dealers use the interdealer market to share the inventory risk they undertake in the public trading round. The price that a dealer quotes to a customer in the first stage depends on the risk-sharing opportunity in the second stage.

4 Simulation Results

4.1 Parameters of trading environments and customer attributes

To determine a customer's preferred market structure, we compare the expected utility that the customer achieves under two market structures, and across trading environments and customer characteristics. The various trading environments we consider are defined over a set of initial parameter values.

Before examining the results of our experiment, we discuss the motivation behind the selection of each parameter value given in Table 2. The third column in Table 2 lists the parameter values for the baseline or benchmark trading environment, and the parameter values in the last column are used to generate modified trading environments. Note that, in generating a modified trading environment (or changing customer attributes), we adjust only one parameter at a time, while keeping the other parameters at their benchmark values. As noted in Table 2, the realized liquidation

Table 2: Parameter Configuration

Panel A: Trading Environment Configuration			
Parameters	Definition	Baseline	Modified
σ_F^2	Variance of liquidation value	0.7	
γ	Risk-aversion differential ($\lambda_c - \lambda_D$)	0	0.5
N	Number of dealers	6	12
$\sigma_{X^D}^2$	Variance of dealer inventory	0	0.7
Panel B: Customer Attributes			
Parameters	Definition	Baseline	Modified
$(\sigma_{X^c}^2, \rho)$	Variance and correlation of inventory shock of customers	(0.7,0)	(1,0) (0.7,0.3)

value of the security is assumed to be drawn from a normal distribution with mean zero and variance, σ_F^2, equal to 0.7 throughout our analysis.

γ defines the relative difference between the customer's level of risk aversion and the dealer's level of risk aversion. In the baseline trading environment, we assume that customers and dealers are equally risk-averse, with the aversion value from their expected utility functions set to 1, so that $\gamma = 0$. That is, in most trading environments considered, $\gamma = 0$. We consider trading environments where γ=0.5, however, implying that dealers are less risk-averse than customers. This trading environment is meant to capture dealers' specialization in market-making, which implies a superior ability to take on inventory risks.

The trading environment is also adjusted by changing the number of market-makers, N. In this way, we can examine the implication of changing market capacity. Specifically, we set the number of market-makers, N, equal to 6 in the baseline configuration and equal to 12 in an alternative configuration of the trading environment.

The degree of dealer heterogeneity can also affect the trading environment (Dutta and Madhavan 1997). We allow dealers to be either homogeneous or heterogeneous in terms of their initial inventory positions prior to trading with customers. Dealer inventory heterogeneity is intended to capture inventory heterogeneity documented by Reiss and Werner (1998). It is also intended to proxy for dealer inventory disturbances incurred during the dealers' public trading in prior periods. Initial dealer positions for dealers $i = 1, ..., N$, are drawn from a normal distribution with mean zero and variance $\sigma^2_{X^D}= 0.7$ when dealers are heterogeneous, and $\sigma^2_{X^D} = 0$ when dealers are assumed to be homogeneous.

To examine how a customer's attributes affect their choice of market structure for a given trading environment, we adjust the characteristics of the liquidity shock that customers receive. To model large versus small customers, we adjust the variance of the distribution from which the liquidity shocks are drawn. Recall that customers are motivated to trade because the liquidity shocks they receive require them to rebalance their portfolios. The larger the liquidity shock, the greater the need for portfolio rebalancing, and the greater the order size that the customer submits to the trading system. We assume that customers who manage large portfolios are more likely to receive large liquidity shocks. As a result, customer liquidity shocks for customer i=1,..., M, are drawn from a normal distribution with mean zero and variance, σ^2_X = 0.7, when customers are characterized as being small, and $\sigma^2_X = 1$ when customers are assumed to be large.

Although the assumption that the initial inventory shock, X_i, is independent across customers seems to be well-justified in the context of equity trading, this assumption does not appropriately capture the stylized facts observed in fixed-income or foreign exchange, where there is a tendency for customer buy (sell) orders to arrive in bunches. Anecdotal evidence indicates that periods of one-sided order flow occur relatively more frequently in fixed-income and FX markets than in equity markets. These order-flow dynamics likely reflect the arrival of public information about the fundamental value of the security. Given that the data release is observed by all market participants, customers would tend to trade in the same direction to adjust their portfolios. In our one-shot game construct, we proxy this order-bunching effect by allowing liquidity shocks, $X_i, i = 1, 2...M$, to be correlated across the customers in some experiments. Because customer order submission strategies are in part driven by their realized liquidity shocks, making these shocks correlated across customers will induce some degree of correlation in customer order flow. We set the correlation between a customer's liquidity shock to ρ=0.3, to engender correlated order flow. The baseline setting is ρ=0, which implies that liquidity shocks are independent across customers and that order flow is, on average, uncorrelated across customers.

4.2 Welfare results

Tables 3 and 4 report customer welfare for each experiment. Each cell in these tables represents one experiment based on 1000 realizations for each normally distributed parameter. Customer welfare reported in each cell is the average across M customers and, for each customer, the mean of the utility level achieved over the 1000 draws of the random variables. As discussed in section 4.1, we use a smoothing technique to optimize the discontinuous utility function. Since we used a negative exponential utility function defined in equation (1), smaller negative values in the cells of Tables 3 and 4 represent an improvement in welfare. Figure 5 shows the improvement in one typical customer's welfare through trading. The lower curve shows their reservation utility in the case of no trading, the black dots show the realization of their utility at the given shocks, and the upper curve shows the fitted values for all the realizations of the utility using a six-degree polynomial. Figures 1 and 2 show the results using cubic spline interpolation on the data in Tables 3 and 4.

Each row of Tables 3 and 4 corresponds to a different trading environment, one where that environment changes with the given number of customers, M. The second and third columns in Table 3 correspond to welfare in the baseline configurations of the order-book and the dealership market structures, respectively. All modified trading environments are compared with those benchmark results.

Table 3 shows that, in the baseline configuration, increasing the number of customers in the market implies improved customer welfare in the order-book system. In an order-book market structure, increasing the number of customers increases the competition for order flow, which reduces the bid-ask spread, and increases the probability of the execution. (We discuss these measures later in this section; see Tables 5 and 6.) An alternative way to think about this is to recall that small values of M imply a trading environment characterized by a low order-arrival rate. Given that customers seek to, in essence, share their inventory risks (i.e., diversify their portfolio imbalances) by trading among a larger set of customers, smaller values of M imply a reduction of potential counterparties in the order-book setting and reduce the avenues for risk diversification for customers, lowering their expected utility.[11]

In contrast, in a dealership market structure with a fixed number of market-makers, an increase in the number of customers leads to lower customer welfare; see column 3 of Table 3. An increase in the number of customers leads to an increase

[11] Vogler (1997) and Rustichini, Satterthwaite, and Williams (1994) show that, as the number of players increases, the double-auction outcome quickly converges to the standard Walrasian equilibrium (i.e., the competitive equilibrium), which is the market structure that provides the maximum level of welfare for customers.

Table 3: Customer Welfare across Trading Environment

Number of customers	Order-book market	Dealership market			
	Baseline	Baseline	Dealer heterogeneity	Positive γ	Number of dealers
M=4	-1.1482	-1.0585	-1.0325	-1.0311	-1.0557
M=6	-1.0936	**-1.1017**	-1.0717	-1.0480	-1.0599
M=9	-1.0700	-1.1039	-1.0960	-1.0520	-1.0660
M=12	-1.0561	-1.1107	-1.1100	-1.0533	-1.0660
M=18	-1.0442	-1.1194	-1.1182	-1.0544	-1.0687
M=24	-1.0364	-1.1335	-1.1314	-1.0585	-1.0818

in the inventory risk faced by dealers. Dealers, in turn, charge wider bid-ask spreads to compensate for the greater inventory risk that they have to take on. As a result, customer welfare falls. The risk-sharing feature of this market is shown in Figure 6, where, after trading, a dealer's final position (black line) shifts away from their initial position (red line) towards the zero line. This implies that dealers share their inventory risk by trading among themselves and that the risk-sharing capacity of the market can be characterized by the slope of the final-position curve. The flatter the line the greater the risk-sharing capacity.

The pivotal parameter is the number of customers. M^*, the critical value of M, is the minimum number of customers required in a market such that their welfare achieved from trading in the order-book systems exceeds their welfare from trading in the dealership system. In this regard, the pivotal value of welfare corresponds to

M^*=6 for the baseline configuration. That is, in the baseline configuration (defined in Table 2), when the number of customers reaches 6 or more, customer welfare is greater in the order-book system than in the dealership system.

The particular value of M^* does not have any empirical significance, since the parameter values used in the simulations have not been calibrated in relation to empirical benchmarks. How the value of M^* changes across various configurations of the trading environment and customer attributes, however, is important and provides relevant insights.

Columns 4, 5, and 6 in Table 3 report customer welfare levels achieved in the dealership market under modified trading environments. Given that the various trading environments correspond to changes in the dealership trading environment only, we do not update the simulation results for the order-book system, since they are unaffected by these changes.

Column 4 reports the results from a trading environment in which dealers hold heterogeneous inventories prior to entering the market. The critical number of customers in this trading environment rises to between M^*=6 and M^*=9, compared with M^*=6 in the baseline trading environment. (Note: The new M^* is actually 7, according to Figure 1.) This implies that a greater number of participants is required for the order-book system to dominate the dealership market. As well, the higher welfare values in column 4 relative to the baseline dealership structure in column 3 indicate that customers are better off when dealers are heterogeneous in terms of their initial inventories. Heterogeneous dealer inventories increase the liquidity risk-sharing capacity of the dealership structure. Since dealers at the extreme of the inventory distribution post the most competitive quotes, customers with opposite trading desires will be able to trade at better prices than in the baseline trading environment. In the benchmark dealership environment, customers at the extreme of the liquidity shock spectrum face a homogeneous set of dealers quoting identical prices. In effect, in this modified environment, dealers use the first stage of trading as a way to hedge their initial inventory positions by offering better bid-ask spreads to attract customers' orders, which therefore benefits the customers.[12]

Column 5 in Table 3 reports the simulation results from a trading environment in which the dealers are less risk-averse than customers, such that $\gamma = 0.5$. In the baseline trading environment, dealers and customers were assumed to be equally risk-averse. The critical number of customers in this trading environment rises between $M^* = 12$ and $M^* = 18$. This implies that a greater number of participants is required for the order-book market to prevail under this trading environment than under the benchmark environment. Given that a dealer's bid-ask spread is a function of their risk aversion, less-risk-averse dealers quote narrower bid-ask spreads in this trading environment, thus benefiting the customer. This result is consistent with the theoretical predictions found in the market-microstructure literature that considers risk-averse market participants (see O'Hara 1995 and Madhavan 2000).

Column 6 in Table 3 reports customer welfare resulting from a trading environment in which the number of dealers, N, is set to 12. The critical number of customers rises relative to the baseline environment to M^*=12. The experiment yields a larger M^* than the baseline trading environment, which implies that a greater number of

[12] For the dealers in this modified trading environment, the first stage of trading serves the same purpose as the second stage of trading. This would seem to imply that interdealer trading is unnecessary, because dealers can offset their inventory risk in the first stage of trading, and this would seem to question the need to model a dealership market structure as a two-stage trading process. Given the limited number of dealers and customers, however, dealers are not able to completely offset their undesired inventory positions in the first stage of trading with the public. Reiss and Werner (1998) indicate that London Stock Exchange (LSE) dealers, for example, have difficulty managing their inventory positions effectively in the public trading stage, since customer order flow does not necessarily gravitate to the most aggressive posted quote. Reiss and Werner offer the broad use of preferencing arrangements as a reason for this lack of competitive pricing power.

customers is required in this case than in the baseline trading environment for an order-book market to dominate a dealership market. Since more dealers participate in second-stage interdealer trading, which improves the risk-sharing capacity in the market, they can quote narrower bid-ask spreads to customers in the first stage. Therefore, we see an increase in customer welfare compared with the baseline case, where there were six market-makers. Similarly, models of dealership markets that ignore private information have shown that, as the number of dealers increases, greater competition among dealers or better risk-sharing opportunities reduces the bid-ask spreads offered to the public (see, for example, Grossman and Miller 1988 and Vogler 1997).

Concern has recently been expressed by policy-makers about the effects that bank mergers may have on market liquidity. As illustrated here, reducing the number of market-makers from N=12 to N=6 reduces the market-making capacity of the dealership market structure and in turn market liquidity, making customers worse off as reduced market liquidity is reflected in wider bid-ask spreads. However, D'Souza and Lai (2002) show that the effect of bank mergers on market liquidity may be ambiguous, since the newly merged bank, via economies of scope, may have a greater risk-bearing capacity, increasing the market-making capacity for the whole of the dealership market.

Our results illustrate the findings of D'Souza and Lai (2002) when we combine the trading environment represented in columns 5 and 6 of Table 3. Specifically, we investigate what happens to M^* when the number of dealers declines from N=12 to N=6 and at the same time the dealers become less risk-averse. M^* in this modified trading environment is best illustrated in Figure 1. The two dealership curves that lie near the top of the chart depict the results as M is varied. The top dealership curve represents the trading environment where there are six dealers and γ =0.5, and the second curve plots the results for the trading environment where there are 12 dealers and γ=0. As the number of dealers declines and their risk aversion declines, M^* rises from 10 to about 13, indicating that the dealership system is more likely to prevail even as the number of dealers declines. In this set-up, the decline in dealer risk aversion more than offsets the negative impact on the risk-sharing capacity of dealership markets arising from a decline in the number of dealers. As such, our framework predicts that investment bank mergers do not unambiguously imply a reduction in the market quality of dealership markets.

Table 4 reports the welfare values from simulations where customers are subject to large (high-variance) liquidity shocks. Large liquidity shocks affect the results for both the order-book and dealership market structure. First, the critical number of customers rises to $M^* = 9$ from $M^* = 6$ in the baseline setting (Panel B). This implies that an increased number of customers is required for an order-book market to prevail when customers are subject to large liquidity shocks. Figure 2 illustrates the results from Table 4 based on a cubic spline interpolation. Note that customer welfare decreases in both the order-book and dealership market structures relative to the baseline environment, but the dealership welfare values are less affected than the order-book system values when M is relatively small.

The results can be explained as follows. In the dealership market structure, dealers are modelled as posting bid-ask prices but are assumed to accept orders of any size at those prices. The dealer's mandate of accepting any order size, C_j, at their posted bid-ask prices is more beneficial for customers who supply large orders to the market than it is for customers who are relatively homogeneous and are subject to small liquidity shocks. Moreover, these customers with large orders benefit from second-stage interdealer trading in the dealership market structure. Dealers compete for large orders that they can then lay off with other dealers in the interdealer setting. In essence, the customers subject to large, high-variance shocks benefit from the extra risk-sharing capacity offered by the dealership market structure.

Table 4: Altering Customer Attributes: Impact of Large and Correlated Shocks

Panel A: Baseline environments		
Number of customers	Order-book market	Dealership market
M=4	-1.1482	-1.0585
M=6	-1.0936	-1.1017
M=9	-1.0700	-1.1039
M=12	-1.0561	-1.1107
M=18	-1.0442	-1.1194
Panel B: High variance of customer liquidity shocks		
Number of customers	Order-book market	Dealership market
M=4	-1.2549	-1.1365
M=6	-1.1783	-1.1428
M=9	-1.1368	-1.1662
M=12	-1.1146	-1.1816
M=18	-1.0985	-1.2041
Panel C: Correlated customer liquidity shocks		
Number of customers	Order-book market	Dealership market
M=4	-1.2787	-1.1764
M=6	-1.2448	-1.2007
M=9	-1.2229	-1.2169
M=12	-1.2224	-1.2484
M=18	-1.2107	-1.2698

When customers are characterized by correlated liquidity shocks (Panel C), we find that again the dealership market is more likely to prevail than in the baseline setting. Relative to the baseline environment, the critical number of customers rises to $M^* = 12$. From Table 4, it is clear that environments characterized by correlated order flow challenge the inventory risk-management capabilities of both market structures, as average customer welfare declines relative to the baseline cases. Again, however, the dealership market structure handles this better than the order-book market structure, which implies a greater M^* value than in the baseline case. In sum, when customers are subject to large, correlated liquidity shocks, dealership markets tend to dominate the order-book system.

4.3 Measures of execution quality

Previous empirical research comparing market structures has, in general, used a single measure of execution quality, the bid-ask spread, as a metric in evaluating the viability of the market structures. Given the multidimensional nature of execution quality, it is important that intermarket comparisons be based on either a broad set of execution quality measures or on some measure directly related to customer welfare. Although our approach provides results on customer welfare across market structures, we can also calculate common measures of execution quality. In this section, we therefore examine in greater detail the link between customer welfare and execution quality. Specifically, for the experiments in which we vary customer attributes, we calculate the bid-ask spreads observed in the order-book system and the public segment of the dealership system, and the probability of execution observed in the order-book system. Since dealership systems ensure the execution of all public orders under all trading environments, the probability of execution is a measure that is pertinent only for limit-order-book systems. Customer welfare is negatively related to bid-ask spreads and is positively related to the probability of execution.

Table 5 gives the bid-ask spreads calculated in three different cases: the baseline environment, when customers are subject to large (more volatile) liquidity shocks, and when customers are subject to correlated liquidity shocks. The bid-ask spreads reported in each cell of Panel A in Table 5 represent the average bid-ask spread submitted to the trading system by the set of M customers. Each cell in Panel B reports the average bid-ask spread posted by dealers in the public segment of the dealership system.

In the order-book market (Panel A), for a given number of M, bid-ask spreads are larger than those observed in the baseline case when customers are subject to higher-variance liquidity shocks. For the dealership market (Panel B), the rise in bid-ask spreads is less substantial, in general. When customers are subject to correlated liquidity shocks, however, the increase in bid-ask spreads relative to the baseline case is greater in the order-book system and is substantially greater in the dealership market. As liquidity shocks increase in size or become correlated, the orders submitted by customers increasingly tax both the order-book and dealership market's ability to provide rebalancing services, and thus bid-ask spreads increase to compensate market participants that are more likely to be left with undesired portfolio or inventory positions.

Table 6 reports the probability of execution on an order-book market. Each number is calulated as following $(1/M)(1/T)\sum_{i=1}^{M}\sum_{t=1}^{T}(I_{(B_{it}>P_t)}+I_{(A_{it}<P_t)})$. When customer liquidity shocks are correlated, the simulation results reported in Table 6 indicate that the average probability of limit-order execution (across customers) decreases from those observed in the baseline case. The decline in the probability of execution reflects the smaller number of offsetting (counterbalancing) customer orders being submitted to the market when orders are correlated. However, the presence of customers in the order-book market who are subject to larger (more volatile) liquidity

Table 5: Bid-ask Spreads in the Order-book and Dealership Markets

Panel A: Bid-ask spread in order-book market			
Number of customers	Baseline: variance of the shock is 0.7	Variance of the shock is 1.0	Correlation of the shock is 0.3
M=4	0.4977	0.5902	0.4962
M=6	0.4193	0.4873	0.4316
M=9	0.3789	0.4079	0.4282
M=12	0.2911	0.3550	0.4267
M=18	0.2802	0.3347	0.4204
M=24	0.2490	0.2976	0.4194
Panel B: Bid-ask spread in dealership market			
Number of customers	Baseline: variance of the shock is 0.7	Variance of the shock is 1.0	Correlation of the shock is 0.3
M=6	0.3298	0.3718	0.8870
M=9	0.3388	0.4382	1.0035
M=12	0.3612	0.4533	1.0357
M=18	0.4213	0.4832	1.0766
M=24	0.4638	0.5012	1.0978

shocks increases the probability of execution. The probability of execution increases because customers are more heterogeneous in this case, and thus are more likely to be matched with customers submitting offsetting orders. Because the probability of execution in the dealership market structure is always one, any increase in the order-execution probability observed in the order-book market would tend to increase the likelihood of that market prevailing, other things equal.

The effects on the order-book market's execution quality seem ambiguous when one combines these higher-execution probability figures with the wider bid-ask spread measures for large (high-variance) customers, reported in Panel A of Table 5. As Table 4 shows, however, customer welfare is in fact lower than in the baseline environment when the liquidity shock variance increases. The effect on customer welfare from wider bid-ask spreads tends to outweigh the beneficial effect of higher probability of order execution. This illustrates the danger of gauging market quality solely on a single measure. The reliance on one measure of execution quality, in this case the probability of execution, would have led to the wrong conclusion about the impact on customer welfare.

5 Conclusions

In this paper we examined the factors that affect a liquidity-motivated investor's preference of market structure. We supposed that customers could choose between a dealership market and a limit-order-book market. We found that the customer's choice of trading venue would vary depending on the trading environment in which they traded and the customer's own characteristics. This study builds on the theoretical market-microstructure literature in that it has developed a theoretical model for both the dealership and order-book markets. To capture the variety in trading environments and customer characteristics observed in actual financial markets, how-

Table 6: Probability of Execution on an Order-Book Market

Number of customers	Baseline: variance of the shock is 0.7	Variance of the shock is 1.0	Correlation of the shock is 0.3
M=4	0.5328	0.5563	0.5011
M=6	0.5754	0.6503	0.5712
M=9	0.6254	0.6683	0.5734
M=12	0.7317	0.7239	0.5898
M=18	0.7502	0.7511	0.5953
M=24	0.7697	0.7785	0.6023

ever, we use an agent-based computational approach that allows us to examine the equilibrium properties of models that are richer than the extant literature in terms of the heterogeneity of trading and customer attributes.

The main predictions of our model can be summarized as follows. First, an increase in the thickness of the market, measured by the number of customers submitting orders at any point in time, causes a rise in the probability of an order-book system prevailing. Second, customers who submit large-size order flow and customers whose order flow tends to be correlated—such as institutional investors—are more likely to trade on dealership trading systems than customers who are smaller and whose order flow is hardly correlated—characteristics generally found in retail investors. Third, a decrease in the number of dealers increases the likelihood that the order-book market will prevail as the preferred trading venue of risk-averse customers. Fourth, an increase in the degree of risk-averse dealer heterogeneity or lower dealer risk-aversion levels (relative to customer risk aversion) increases the likelihood that customers will choose to trade on the dealership system.

Our findings can be understood in terms of a market structure's capacity to supply liquidity-risk-sharing services under different trading environments, combined with the idea that the varying demand for these services depends on various customer attributes. As trading environments are changed or as customer attributes are varied, a particular market structure will be more robust in terms of its capacity to provide risk-sharing services to market participants and in turn become the preferred trading venue. Overall, our study shows that policy-makers should be leery of adopting regulations that explicitly or implicitly force competing trading venues (i.e., competing trading systems) to adopt identical or similar trading architectures. This type of "one shoe [market structure] fits all" policy would in the end be detrimental to public investors, who are exactly the agents that regulators seek to protect.

The results also help to explain several of the empirical regularities observed in financial markets. First, they explain the appearance of multiple trading venues, each with a different trading structure in various financial market segments. Given the heterogeneity of customer characteristics, different market structures will appeal to different customers. To the extent that some customer types lacked access to their preferred market trading system, the entry of new trading venues that offered trading structures that appealed to this underserved customer would explain their success in competing with incumbent trading venues. Second, our model explains why the public segments of the government bond and FX markets, where trading is driven almost exclusively by institutional investors, are predominantly organized as dealership markets.

Our model also predicts that if the merging of investment banks leads simply to a decline in the number of market-makers in a dealership market, then a competing order-book system is more likely to succeed. If, however, after investment bank mergers, the newly formed market-makers are less risk-averse than the individual pre-merger market-makers, perhaps because of economies of scope, then the impact on the customer's preferred market structure is ambiguous.

Several questions that lie outside of the scope of this study have been left for future research. First, we did not allow for asymmetric information to play a role in the dealer's or the customer's optimizing behaviour. Given that our traders did not act on private information, we cannot examine how the public disclosure of trade information may impact customer welfare and execution quality in various trading environments. Our approach is sufficiently flexible, however, to allow our modelled market participants access to private signals about the liquidation value of the security. We will pursue this research in a companion paper.

Second, we did not study the degree of market efficiency offered by the two market structures examined. A key question raised in much of the agent-based literature is the extent to which the efficiency (in the Pareto sense) of any given market mechanism is attributable to the trading agent's behaviour versus the inherent design of the market structure. It would be interesting to measure the degree of market efficiency achieved in the trading mechanisms examined here and then assess, perhaps by endowing agents with bounded learning mechanisms, the extent to which market efficiency is attributable to these various factors.

Finally, Canadian regulators recently considered rules that would, in effect, have allowed customers in the government securities dealership market to submit orders to the IDB systems rather than trade with dealers. This would have changed the nature of the decision-making process of both the customers and dealers, since customers would have direct access to the inventory risk-sharing services provided by the IDB system rather than indirect access to these services through dealers. Future research on this issue could examine three market trading mechanisms: the two studied in this paper, plus a third, hybrid dealership structure where some customers submit orders to the IDB system while other customers continue to trade with dealers.

6 Appendix A

This appendix describes the numerical procedure used throughout this study to find the optimal strategies for each dealer and the resulting equilibrium. The procedure we have developed can be used to solve a wide class of one-period simultaneous-move games. Under certain conditions, one can also use this method to solve multiperiod games, such as in a dealership market and an order-book system, which are considered in this paper.

Before describing the numerical procedure, we note the following two caveats. First, even though a wide range of games can be solved using our method, we do not seek to define that class of games. Second, if a game admits multiple equilibria, there is no guarantee that the method we propose can solve numerically for all of them. Moreover, depending on their stability characteristics, certain types of equilibria cannot be solved using our method.

Section 6 describes the type of game our method is intended to solve and the method itself. Sections 6.2 and 6.3 describe their application to the order-book market and dealership market.

6.1 Description of the Algorithm

6.1.1 Necessary conditions for the Nash equilibrium

Let Γ be a simultaneous-move game described by the following elements. Let $\mathcal{N} = \{1, \ldots, N\}$ be a set of N players. Each player observes a particular random event, ω_i, drawn from a distribution with support $\Omega_i \subset \mathcal{R}^M$. Each player also has a strategy $S_i : \omega_i \longmapsto a_i \in \mathcal{A}_i$,[13] which maps a random event, ω_i, to an action, a_i ; $\mathcal{A}_i$ is a set of actions. The game also defines for each player a payoff function, $U_i : \left((a_i)_{i\in\mathcal{N}}, \varepsilon\right) \longmapsto$

[13] We restrict ourselves to games where the players have deterministic strategies.

u_i, a real number depending on the realization of a random variable, ε. Therefore, the problem faced by each player is to find a strategy, $s_i^* \in S_i$, such that

$$s_i^*\left((s_j)_{j\in\mathcal{N}^{i-}}\right) = \arg\max_{s_i\in S_i} EU_i\left((s_j(\omega_j))_{j\in\mathcal{N}}, \varepsilon\right), \tag{17}$$

where $\mathcal{N}^{i-} = \{1, ...i-1, i+1, ...N\}$. The function s_i^* will be called the best-response function of player i. In the case where

$$s_i^{**}\left(\left(s_j^{**}\right)_{j\in\mathcal{N}^{i-}}\right) = s_i^{**} \qquad \forall i \in \mathcal{N}, \tag{18}$$

we will call the set of strategies $(s_i^{**})_{i\in\mathcal{N}}$ a Nash equilibrium of the game Γ. As (18) indicates, a Nash equilibrium is a fixed point of the best-response function of all the players taken simultaneously.

This section describes how this fixed-point problem can be solved numerically. As will be explained later, this procedure optimizes numerically the utility function of each player's possible strategies. To guarantee that a numerical optimization is possible, we require that the game we analyze meet the following two conditions:

1. Each strategy function can be parameterized by a real-valued vector of parameters θ_i. We want the strategy function to have a parametric form because the numerical optimization of each player's utility will be substantially simplified if we optimize over a set of parameters rather than a set of functions.

2. The utility function $U_i\left((s_j(\omega_j, \theta_j))_{j\in\mathcal{N}}, \varepsilon\right)$ is continuously differentiable on θ_i. This condition is used to guarantee that the first derivative of the utility function is continuous. If this were not the case, it would be much more difficult to define the optimal response of a player; moreover, it would be more difficult to obtain that response numerically.

If these requirements are met, a necessary condition for a set of strategies $(s_i^{**})_{i\in\mathcal{N}}$

to be a Nash equilibrium is

$$\frac{\partial EU_i}{\partial \theta_i}\left(\left(s_j^{**}\left(\omega_j;\theta_j\right)\right)_{j\in\mathcal{N}^{i-}},\varepsilon\right)=0 \qquad \forall i\in\mathcal{N}. \tag{19}$$

This condition states that no player has an incentive to deviate, at least locally, from s_i^{**}.

Unfortunately, many interesting problems of game theory exist where these two conditions cannot be satisfied. In those cases, the characterization of the Nash equilibrium given in (19) is not valid.

Nevertheless, it will often be possible to modify a game so that both conditions are met. Doing so will allow us to apply our numerical procedure to the modified game. More specifically, the two conditions can be treated as described below.

For the first condition, suppose that we know very little about the form of s^*. It is still possible to use a flexible function that is parameterized by a vector of parameters, θ, that can be arbitrarily close to the real optimal strategy function, s^*. Numerous examples of such functions exist, such as splines or kernel-based functions.

For the games we analyze in this study, we use an ANN that has a form given by

$$\begin{aligned} a_i &= \tilde{s}_i\left(\omega_i;\theta\right) \qquad\qquad (20)\\ &= ANN\left(\omega_i;\theta\right)\\ &= \theta_{i,0}+\sum\nolimits_{k=1}^{H}\theta_{i,k}\ \varphi\left(\theta_{i,(H+1)+k}+\sum\nolimits_{j=1}^{J}\theta_{i,(2H+1)+j}\ \omega_{i,j}\right), \end{aligned}$$

where a_i is the action performed by player i with given information, ω_i, and a strategy described by the vector of parameters, θ. J is the total number of inputs. The function φ is generally chosen as a twice continuously differentiable function from $\mathcal{R}$ to a bounded interval of $\mathcal{R}$, generally $(0,1)$ or $(-1,-1)$. A few standard choices for the activation function φ are tanh, sigmoid, or a bell-shaped Gaussian function shown, respectively, as

$$\begin{aligned}\varphi(x) &= \tanh(x)\\ \varphi(x) &= (1+\exp(-x))^{-1}\\ \varphi(x) &= \exp(-x^2).\end{aligned}$$

Often, this form is described in terms of its number of input nodes, which is equal to the dimensionality of ω_i and its number of hidden nodes, given by the parameter H. This terminology refers to the layered organization of the ANN specification. More explicitly: the value of the argument $\omega_{i,j}$ is referred to as the activation level of the input node j, the value of $\varphi(.)$ is called the activation level of a hidden (or middle) node, and a_i is called the activation level of the output node.

This specification is known to be a general approximator of a very wide class of functions, given that the number of hidden nodes, H, is large enough (White 1992). Therefore, $\tilde{s}_i$ can be used as an approximation of s_i.

For the second condition, because U_i is generally not a once continuously differentiable function of θ_i, we first define an increasing sequence of real numbers $\{\eta_k\}_{k\in\mathcal{R}}$ s.t. $\lim_{k\to\infty}\eta_k=\infty$ and then look for a sequence of functions, $\left(\tilde{U}_i^{\eta_k}\right)_{\eta\in\mathcal{R}}$ that are each continuously differentiable on θ_i and tend to U_i as $k\to\infty$.[14] We will refer to the variable η_k as the smoothing parameter. The particular choice of $\tilde{U}_i^{\eta_k}$ that we made for this study will be explained in section 6.1.2, where we describe more explicitly the application of our procedure to the games we analyze. At this point, one should think of $\tilde{U}_i^{\eta_k}$ as given.

[14] The formal demonstration of the convergence of the smoothing methods is beyond the scope of this study and we leave it for future research.

6.1.2 Numerical procedures for the Nash equilibrium

Once a parametric form of the strategy function and a continuously differentiable utility function are defined, it is possible to numerically find an equilibrium of the modified game using the new strategy and utility functions. The method we use here is as follows: for a given value of η we have a set of functions $\left(\tilde{U}_i^{\eta_0}\right)_{i\in\mathcal{N}}$ that defines the game $\tilde{\Gamma}^{\eta_0}$, which has an equilibrium $\left(\tilde{s}_i^{**\eta_0}\left(\omega_i;\theta_i\right)\right)_{i\in\mathcal{N}}$ characterized by (19). To get the equilibrium for Γ, we will first set a relatively low value η_0 for the smoothing parameter and search the equilibrium strategies $\left(\tilde{s}_i^{**\eta_0}\left(\omega_i;\theta_i\right)\right)_{i\in\mathcal{N}}$. Then we increase the smoothing parameter to η_1, and track the movement of the equilibrium from $\left(\tilde{s}_i^{**\eta_0}\left(\omega_i;\theta_i\right)\right)_{i\in\mathcal{N}}$ to $\left(\tilde{s}_i^{**\eta_1}\left(\omega_i;\theta_i\right)\right)_{i\in\mathcal{N}}$. The procedure is continued using a sequence $(\eta_2,\eta_3,...)$ that tends to infinity. This procedure yields a sequence of equilibria $\left(\left(\tilde{s}_i^{**\eta_k}\left(\omega_i;\theta_i\right)\right)_{i\in\mathcal{N}}\right)_{k=1,2,...K}$ of the sequence of games $\left(\tilde{\Gamma}^{\eta_k}\right)_{k=1,2,...K}$ that tends to the equilibrium of Γ.

More explicitly, for a given η_k, it is possible to approximate numerically the equilibrium of the game $\tilde{\Gamma}^{\eta_k}$ by using the fixed-point characterization given in (18) or (19). The procedure is as follows:

Step 1: Draw a large number, T, of realizations of the random variables $\left(\omega_i\right)_{i\in\mathcal{N}}$, and ε should be drawn from the appropriate distribution. These variables will be denoted $\left(\omega_i^t\right)_{i\in\mathcal{N}}$ and ε^t, $t=1,....,T$.

Step 2: For $i=1$, and given $\left(\theta_j\right)_{j\in\mathcal{N}}$, we numerically search for

$$\arg\max_{\theta_i}\frac{1}{T}\sum_{t=1}^{T}\widetilde{U}_i^{\eta_k}\left(\left(s_j\left(\omega_j^t;\theta_j\right)\right)_{j\in\mathcal{N}},\varepsilon^t\right), \tag{21}$$

and assign that value to θ_i. This procedure corresponds to finding the optimal response of player i to the strategies $\left(s_j\left(\omega_j^t;\theta_j\right)\right)_{j\in\mathcal{N}}$. This is the equivalent to solving the equations given in (19) for θ.

Step 3: The variable i is increased incrementally to $i+1$, and step 2 is repeated until convergence of all the $\left(s_j\left(\omega_j^t;\theta_j\right)\right)_{j\in\mathcal{N}}$.

After applying this procedure, an equilibrium is obtained for the game $\tilde{\Gamma}^{\eta_k}$. The smoothing parameter is then increased from η_k to η_{k+1}, and we search for an equilibrium of the game $\tilde{\Gamma}^{\eta_{k+1}}$ using as starting values for θ the values obtained at the equilibrium of $\tilde{\Gamma}^{\eta_k}$. Repeating this procedure until η is large enough, we obtain the equilibrium of a game that is an arbitrarily close approximation of the game Γ.

The method we use to find the Nash equilibria is similar to searching for a fixed point in a function $f(x)$ by starting at point x_0 and iterating on $x_{i+1} = f(x_i)$. It is well known that this method will converge to a fixed point, $x^*$of f, given that, starting for x_0 close enough to x^*, the function f must respect $\left|\frac{\partial f}{\partial x}(x^*)\right| < 1$. In our case, the analogue of this condition can be understood as a stability condition on the equilibrium that is similar to the idea of the trembling-hand equilibrium: given that a player diverges locally from their equilibrium strategy, iterating on the optimal reaction from all the other players should bring back all those players' reactions towards the equilibrium.

6.2 Application A: Order-Book Market – A One-Stage Game

In an order-book market, we have a set $\mathcal{M}$ of M customers. Each of the customers receives an inventory shock, X_i, and participates in a double auction by posting a bid quote $\left(B_i, Q_i^B\right)$, indicating that if the bid, B_i, is greater than the price, P, the customer agrees to buy Q_i^B units of the risky asset at the price P. Also, the customer will post an ask quote $\left(A_i, Q_i^A\right)$, indicating that if the price is above their ask, A_i, they agree to sell Q_i^A units at the price P. This price in a k-double auction (k-DA) is given by equation (6).

After the auction, the value of the risky asset, F, is revealed. The random variables F and $(X_i)_{i\in\mathcal{M}}$ are independent normal variables with mean 0 and variance σ_F^2 and

σ_X^2, respectively.

The information used by customer i is $\omega_i = (X_i)$. Therefore, the strategy function s_i is a mapping from X_i to $\left(B_i, Q_i^B, A_i, Q_i^A\right)$, where $B_i \leq A_i$ and $Q_i^B, Q_i^A \geq 0$. We can define the wealth of player i by

$$W_i = X_i F + Q_i (F - P), \tag{22}$$

where Q_i is the quantity executed of player i's order and the payoff to player i is in turn given by

$$U_i = -\exp(-\lambda_i W_i), \tag{23}$$

where λ_i is the risk aversion of customer i.

6.2.1 Parameterization of the strategy function

As stated in section 6.1.1, there are two necessary conditions for finding the Nash equilibrium. We first must parameterize the strategy function. In this application we use the following specification for the market participants' strategy functions:

$$\begin{aligned} \frac{B_i + A_i}{2} &= ANN\left(X_i, \theta_i^{mid}\right) \\ A_i - B_i &= \phi\left(\theta_i^{spread}\right) \\ Q_i^B &= \phi\left(ANN\left(X_i, \theta_i^{Q^B}\right)\right) \\ Q_i^A &= \phi\left(ANN\left(X_i, \theta_i^{Q^A}\right)\right), \end{aligned} \tag{24}$$

and therefore θ_i is $\left(\theta_i^{mid}, \theta_i^{spread}, \theta_i^{Q^B}, \theta_i^{Q^A}\right)'$. ANN is defined by equation (20). The function ϕ is given by

$$\phi(x) = \log(1 + \exp(x)),$$

and is used to give non-negative actions.

6.2.2 Modification of the utility function

It can easily be seen that the function U_i does not satisfy the smoothness conditions described earlier. To make the utility function differentiable, a smoothing technique is required to modify it. Note that the smoothing method is not unique as long as the convergence of the utility function at the limit is guaranteed. In this application, we show one example of the smoothing method that can be used in this game.

Given a smoothing parameter, η, we will rewrite the definition of P, now denoted $\tilde{P}^\eta$, as the root of

$$\sum\nolimits_{i=1}^{M} Q_i^B \, \tilde{I}^\eta_{\left(B_i - \tilde{P}^\eta\right)} - \sum\nolimits_{i=1}^{M} Q_i^A \, \tilde{I}^\eta_{\left(\tilde{P}^\eta - A_i\right)} = 0, \tag{25}$$

where $\tilde{I}_\eta$ is defined as

$$\tilde{I}^\eta_{(x)} = \begin{cases} 0 & if \quad x < -\dfrac{1}{\eta} \\ \dfrac{(x\,\eta)^2}{2} + x\,\eta + \dfrac{1}{2} & if \quad -\dfrac{1}{\eta} < x < 0 \\ -\dfrac{(x\,\eta)^2}{2} + x\,\eta + \dfrac{1}{2} & if \quad 0 < x < \dfrac{1}{\eta} \\ 1 & if \quad x > \dfrac{1}{\eta} \end{cases} . \tag{26}$$

Hence, $\tilde{I}_\eta$ is a second-degree polynomial defined by parts having continuous first derivatives. In the limit, as $\eta \to \infty$, it is equivalent to $I\,(x > 0)$. It has the advantage that the space of polynomials defined by parts is closed under addition, implying that (25) is also a second-degree polynomial defined by parts having continuous first derivatives. Therefore, its root can easily be calculated.

The use of function $\tilde{I}_\eta$ can be illustrated in an excess demand function in Figure 3. As the smoothing parameter, η, increases from 1 to 200, this smoothed excess demand function becomes steeper and approaches the step function.

Given that this polynomial indicator function is closed under addition, implying that the excess demand function (25) is also a second-degree polynomial defined by parts, its root, $\tilde{P}_\eta$, can be calculated. Now, given $\tilde{P}_\eta$, the executed quantity for the player, i, is

$$\tilde{Q}_i^\eta = Q_i^B \, \tilde{I}^\eta_{\left(B_i - \tilde{P}^\eta\right)} - Q_i^A \, \tilde{I}_{\left(\tilde{P}^\eta - A_i\right)}, \tag{27}$$

and the wealth of player i in equation (2) becomes

$$\tilde{W}_i^\eta = X_i \, F + \tilde{Q}_i^\eta \left(F - \tilde{P}^\eta\right). \tag{28}$$

Accordingly, the utility function in equation (1) becomes

$$\tilde{U}_i^\eta = -\exp\left(-\lambda_i \, \tilde{W}_i^\eta\right). \tag{29}$$

The objective function of each player is approximated by simulation. We begin this method by drawing T realizations of the vector $(X_i)_{i \in \mathcal{M}}$. Given (X_{it}, θ_i), each player submits a bid-and-ask quote $\left(B_{it}, Q_{it}^B, A_{it}, Q_{it}^A\right)$, from which we can find $\left(\tilde{P}_t^\eta, \tilde{Q}_{it}^\eta\right)$. Conditional on $(X_{jt})_{j \in \mathcal{M}}$, the wealth of player i has a normal distribution with mean $-\tilde{Q}_{it}^\eta \, \tilde{P}_t^\eta$ and variance $\sigma_F^2 \left(X_{it} + \tilde{Q}_{it}^\eta\right)^2$. Therefore, the utility of player i conditional on $(X_{it})_{i \in \mathcal{M}}$ can be written as

$$E\left(\tilde{U}_i^\eta \mid (X_{jt})_{j \in \mathcal{M}}\right) = -\exp\left(\frac{\lambda_i^2 \, \sigma_F^2}{2} \left(X_{it} + \tilde{Q}_{it}^\eta\right)^2 + \lambda_i \, \tilde{Q}_{it}^\eta \, \tilde{P}_t^\eta\right).$$

The objective function for player i is approximated by

$$E\tilde{U}_i^\eta \cong \frac{1}{T} \sum_{t=1}^{T} E\left(\tilde{U}_i^\eta \mid (X_{jt})_{j \in \mathcal{M}}\right). \tag{30}$$

For our experiments, we choose $\sigma_F^2 = 0.7$, $\sigma_X^2 = 0.7$, and $\lambda_i = 1$. These values guarantee that in the case where no trade happens in the double auction, the mean and the variance of the unconditional expected utility are still finite. Also, we set $T = 1000$ for the number of random draws. The smoothing parameter, η, starts with an initial value of 1 and is multiplied by 1.2 on each iteration, until it reaches a value of approximately 1000. In ANN, the structure we choose is one input, one hidden node, and one output. The activation function, φ, takes the form of $\varphi(x) = \tanh(x)$.

6.3 Application B: Dealership Market – A Two-Stage Game

Recall that, in a dealership market, we have a set $\mathcal{M}$ of M customers indicated by j and a set $\mathcal{N}$ of N dealers indicated by i. Each customer receives an inventory shock, X_j. Each dealer, i, will submit a pair of quotes $\left(B_i^1, A_i^1\right)$ indicating that they agree to buy or sell any number of units of the risky asset at a price of B_i^1 or A_i^1, respectively. Each customer, j, will then be allowed to observe a randomly chosen set of k dealers, Δ_j, and decide the quantity, C_{ij}, that they wish to trade, and with which dealer. The total amount traded by customer j is $C_j = \sum_{i\in\mathcal{N}} C_{ij}$. The dealers are then allowed to trade among each other by participating in a double auction, where they submit a bid quote $\left(B_i^2, Q_i^{B2}\right)$ and an ask quote $\left(A_i^2, Q_i^{A2}\right)$. As in an order-book market, a quantity, Q_i, will be executed at a price, P, and a fundamental value, F, will be revealed.

In the first stage of the game, each customer observes a set of quotes $\left(B_i^1, A_i^1\right)_{i\in\Delta_j}$. The maximum bid of that subset, $\overline{B}_j^1$, and the minimum ask, $\underline{A}_j^1$, specifically, are

$$
\begin{aligned}
\overline{B}_j^1 &= \max\{B_i^1 : i \in \Delta_i\} \\
\underline{A}_j^1 &= \min\{A_i^1 : i \in \Delta_i\}.
\end{aligned}
$$

Similarly, as in an order-book market, the values for P and Q_i can be smoothed according to equations (25), (26), and (27).

On the other hand, C_{ij} still has to be written as a smooth function of B^1 and A^1. First, we rewrite the best quotes, $\overline{B}_j^1$ and $\underline{A}_j^1$, as

$$
\begin{aligned}
\widetilde{\overline{B}}_j^{1\eta} &= \frac{\sum_{i\in\Delta_j} B_i^1 \, \exp\left(\eta\, B_i^1\right)}{\sum_{i\in\Delta_j} \exp\left(\eta\, B_i^1\right)} \\
\widetilde{\underline{A}}_j^{1\eta} &= \frac{\sum_{i\in\Delta_j} A_i^1 \, \exp\left(-\eta\, A_i^1\right)}{\sum_{i\in\Delta_j} \exp\left(-\eta\, A_i^1\right)}.
\end{aligned}
\tag{31}
$$

These functions are weighted averages of the B_i^1 and A_i^1 and converge towards the maximum and minimum elements of $\left(B_i^1\right)_{i\in\Delta_j}$ and $\left(A_i^1\right)_{i\in\Delta_j}$, respectively.

Following this, the customer's order, C_j, in function (11) can be rewritten as

$$\widetilde{C}_j^\eta = \left(\frac{-\widetilde{\overline{B}}_j^{1\eta}}{\lambda_j^C\,\sigma_F^2} - X_j\right)\,\widetilde{I}^\eta_{\left(\widetilde{\overline{B}}_j^{1\eta} - P_j^R\right)} + \left(\frac{-\widetilde{\underline{A}}_j^{1\eta}}{\lambda_j^C\,\sigma_F^2} - X_j\right)\,\widetilde{I}^\eta_{\left(P_j^R - \widetilde{\underline{A}}_j^{1\eta}\right)}.$$

The share of C_j in function (13) that will be passed to dealer i is given by

$$\widetilde{C}_{ij}^\eta = \widetilde{C}_j^\eta\left(\widetilde{I}^\eta_{\left(\widetilde{\overline{B}}_j^{1\eta} - P_j^R\right)}\frac{\exp\left(\eta\,B_i^1\right)}{\sum_{i\in\Delta_j}\exp\left(\eta\,B_i^1\right)} + \widetilde{I}^\eta_{\left(P_j^R - \widetilde{\underline{A}}_j^{1\eta}\right)}\frac{\exp\left(-\eta\,A_i^1\right)}{\sum_{i\in\Delta_j}\exp\left(-\eta\,A_i^1\right)}\right). \quad (32)$$

Finally, a smoothed version of dealer i's wealth (15) can be written as

$$\widetilde{W}_i^{D\eta} = \sum_{j\in\mathcal{M}}\widetilde{C}_{ij}^\eta\left(\widetilde{I}^\eta_{\left(\widetilde{C}_{ij}^\eta>0\right)}A_i^1 + \widetilde{I}^\eta_{\left(\widetilde{C}_{ij}^\eta<0\right)}\,B_i^1\right) - \tilde{Q}_i^\eta\,\tilde{P}^\eta + F\left(\tilde{Q}_i^\eta - \sum_{j\in\mathcal{M}}\widetilde{C}_{ij}^\eta\right). \quad (33)$$

The objective function of each dealer is approximated by simulation. We begin this method by drawing T realizations of the vector $\left(X_j\right)_{j\in\mathcal{M}}$ and $\left(\Delta_j\right)_{j\in\mathcal{M}}$. Conditionally on $\left(X_j\right)_{j\in\mathcal{M}}$ and $\left(\Delta_j\right)_{j\in\mathcal{M}}$, dealer i's wealth, $\widetilde{W}_i^{D\eta}$, in equation (33) has a normal distribution with mean and variance as follows:

$$\begin{aligned}
E(\widetilde{W}_{it}^{D\eta}|\left(X_{jt}\right)_{j\in\mathcal{M}},\left(\Delta_{jt}\right)_{j\in\mathcal{M}}) &= \sum_{j\in\mathcal{M}}\widetilde{C}_{ijt}^\eta\left(\widetilde{I}^\eta_{\left(\widetilde{C}_{ijt}^\eta>0\right)}A_{it}^1 + \widetilde{I}^\eta_{\left(\widetilde{C}_{ijt}^\eta<0\right)}\,B_{it}^1\right) - \tilde{Q}_{it}^\eta\,\tilde{P}_t^\eta \\
var(\widetilde{W}_{it}^{D\eta}|\left(X_{jt}\right)_{j\in\mathcal{M}},\left(\Delta_{jt}\right)_{j\in\mathcal{M}}) &= \sigma_F^2\left(\tilde{Q}_{it}^\eta - \sum_{j\in\mathcal{M}}\widetilde{C}_{ijt}^\eta\right)^2.
\end{aligned}$$

Therefore, the utility of dealer i, conditional on $\left(X_{jt}\right)_{j\in\mathcal{M}}$ and $\left(\Delta_{jt}\right)_{j\in\mathcal{M}}$, can be written as

$$\begin{aligned}
E(\widetilde{U}_{it}^{D\eta}|\left(X_{jt}\right)_{j\in\mathcal{M}},\left(\Delta_{jt}\right)_{j\in\mathcal{M}}) &= -\exp[-\lambda^D(E(\widetilde{W}_{it}^{D\eta}|\left(X_{jt}\right)_{j\in\mathcal{M}},\left(\Delta_{jt}\right)_{j\in\mathcal{M}})) \\
&\quad + \frac{(\lambda^D)^2}{2}var(\widetilde{W}_{it}^{D\eta}|\left(X_{jt}\right)_{j\in\mathcal{M}},\left(\Delta_{jt}\right)_{j\in\mathcal{M}})]. \quad (34)
\end{aligned}$$

Using a method identical to the one used in the order-book market, the unconditional expected utility can be approximated by

$$E(\widetilde{U}_i^{D\eta}) \cong \frac{1}{T}\sum_{t=1}^{T} E(\widetilde{U}_{it}^{D\eta} |\, (X_{jt})_{j\in\mathcal{M}}, (\Delta_{jt})_{j\in\mathcal{M}}). \tag{35}$$

For our baseline experiments, we chose a setup similar to the one we used for the order-book system to facilitate the comparison: $\sigma_F^2 = 0.7$, $\sigma_X^2 = 0.7$, $\lambda^D = 1, \lambda^C = 1$. We used $T = 1000$ random draws. The smoothing parameter starts with an initial value of 1 and is multiplied by 1.2 on each iteration, until it reaches a value of approximately 1000.

References

Bank for International Settlements. 2001. "The Implications of Electronic Trading in Financial Markets." CGFS Publication No. 16 (January).

Bessembinder, H. 1999. "Trade Execution Costs on Nasdaq and the NYSE: A Post-Reform Comparison." *Journal of Financial and Quantitative Analysis* 34: 387-407.

Bloomfield, R. and M. O'Hara. 1999. "Market Transparency: Who Wins and Who Loses?" *Review of Financial Studies* 12: 5-35.

—. 2000. "Can Transparent Markets Survive?" *Journal of Financial Economics* 55: 425-59.

Brock, W. 2000. "Whither Nonlinear?" *Journal of Economic Dynamics and Control* 24: 663-78.

Conrad, J., K. Johnson, and S. Wahal. 2002. "Institutional Trading and Alternative Trading Systems." *Journal of Financial Economics*. Forthcoming.

Dechert, W.C. and C. Hommes. 2000. "Complex Nonlinear Dynamics and Computational Methods." *Journal of Economic Dynamics and Control* 24: 651-62.

de Jong, F., T. Nijman, and A. Roel. 1995. "A Comparison of the Cost of Trading French Stocks on the Paris Bourse and on SEAQ International." *European Economic Review* 39: 1277-1301.

D'Souza, C. and A. Lai. 2002. "The Effects of Bank Consolidation on Risk Capital

Allocation and Market Liquidity." Bank of Canada Working Paper No. 2002-5.

Dutta, P. and A. Madhavan. 1997. "Competition and Collusion in Dealer Markets." *Journal of Fiance* 52(1): 245-76.

Financial Services Authority. 2000. "The FSA's Approach to Regulation of Market Infrastructure." FSA Discussion Paper (January).

Flood, M., R. Huisman, K. Koedijk, and R. Mahieu. 1999. "Quote Disclosure and Price Discovery in Multiple-Dealer Financial Markets." *Review of Financial Studies* 12: 37-59.

Ganley, J., A. Holland, V. Saporta, and A. Vila. 1998. "Transparency and the Design of Securities Markets." *Bank of England Financial Stability Review* No. 4, Spring.

Gravelle, T. 2002. "The Market Microstructure of Multiple-Dealer Equity and Government Securities Markets: How They Differ." Bank of Canada Working Paper No. 2002-9.

Grossman, S. and M. Miller. 1988. "Liquidity and Market Structure." *Journal of Finance* 43(3): 617-37.

Keim, D. and A. Madhavan. 1997. "Execution Costs and Investment Style: An Inter-Exchange Analysis of Institutional Equity Trades." *Journal of Financial Economics* 46: 265-92.

LeBaron, B. 2000. "Agent Based Computational Finance: Suggested Readings and Early Research." *Journal of Economic Dynamics and Control* 24: 679-702.

LeBaron, B. 2001. "A Builder's Guide to Agent Based Financial Markets." *Quantitative Finance* 1(2): 254-61.

—. 2002. "Building Financial Markets with Artificial Agents: Desired Goals, and Present Techniques." In *Computational Markets*, Edited by G. Karakoulas, MIT Press. Forthcoming.

Lyons, R. 1996. "Optimal Transparency in a Dealer Market with an Application to the Foreign Exchange Market." *Journal of Financial Intermediation* 5: 225-54.

—. 2001. *The Market Structure Approach to Exchange Rates.* MIT Press.

Madhavan, A. 2000. "Market Microstructure: A Survey." *Journal of Financial Markets* 3: 205-58.

Madhavan, A., D. Porter, and D. Weaver. 2001. "Should Securities Markets be Transparent?" In *Financial Market Structure and Dynamics*, 57-62. Proceedings of a conference held by the Bank of Canada, November 2001. Ottawa: Bank of Canada.

Naik, N., A. Neuberger, and S. Viswanathan. 1999. "Trade Disclosure Regulation in Markets with Negotiated Trades." *Review of Fiancial Studies* 12(4): 873-900.

Nicolaisen, J., V. Petrov, and L. Tesfatsion. 2001. "Market Power and Efficiency in a Computational Electricity Market with Discriminatory Double-Auction Pricing." *IEEE Transactions on Evolutionary Computations* 5(5).

O'Hara, M. 1995. *Market Microstructure Theory.* Oxford: Basil Blackwell.

Reiss, P. and I. Werner. 1998. "Does Risk-Sharing Motivate Interdealer Trading?"

Journal of Finance 53(6): 1657-1703.

Rustichini, A., M. Satterthwaite, and S. Williams. 1994. "Convergence to Efficiency in a Simple Market with Incomplete Information." *Econometrica* 62: 1041-63.

Saporta, V. 1997. "Which Inter-Dealer Market Prevails? An Analysis of Inter-dealer Trading in Opaque Markets." Bank of England Working Paper No. 59 (March).

Securities and Exchange Commission. 2000. "Electronic Communication Networks and After-Hours Trading."

Sirri, E. 2000. "Innovation Around Stasis: The Exchange Market's Response to SEC Regulation of Institutional Form." AEI Working Paper (February).

Smith, B.F., D. Alasdair, S. Turnbull, and R.W. White. 2000. "Upstairs Market for Principal and Agency Trades: Analysis of Adverse Information and Price Effects." *Journal of Finance*. Forthcoming.

Theissen, E. 2000. "Market Structure, Informational Efficiency and Liquidity: An Experimental Comparison of Auction and Dealer Markets." *Journal of Financial Markets* 3: 333-63.

Viswanathan, S. and J.J.D. Wang. 2002. "Market Architecture: Limit-Order Books versus Dealership Markets." *Journal of Financial Markets* 5: 127-68.

Vogler, K.-H. 1997. "Risk Allocation and Inter-Dealer Trading." *European Economic Review* 41: 1615-34.

Werner, I. 1997. "A Double Auction Model of Interdealer Trading." Stanford University Graduate School of Business Research Paper No. 1454 (September).

White, H. 1992. *Artificial Neural Networks: Approximation and Learning Theory.* Cambridge, MA.: Blackwell Publishers.

Yang, J. 2002. "The Efficiency of an Artificial Double Auction Stock Market with Neural Learning Agents." In *Evolutionary Computation in Economics and Finance* edited by S.H. Chen, 85-105. Springer-Verlag.

Figure 1: **Customer Welfare across Trading Environment**

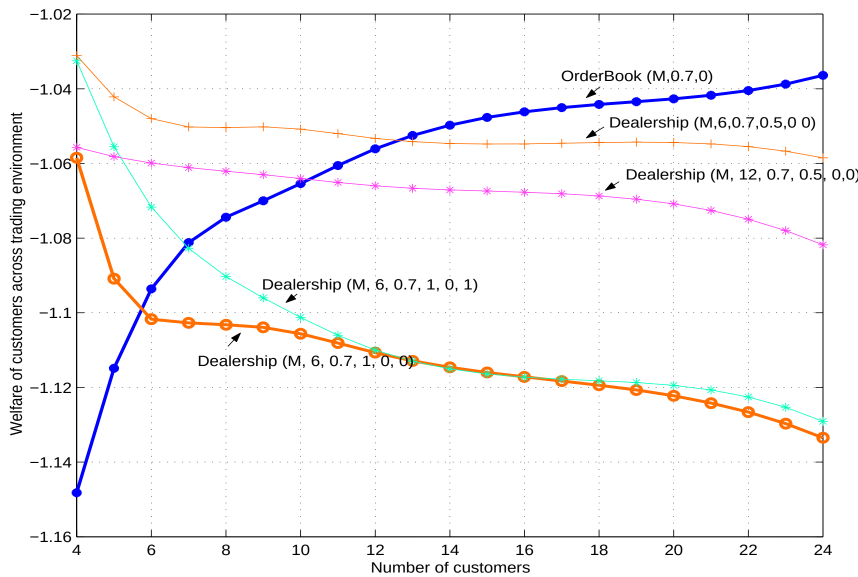

Figure 2: **Impact of the Feature of Customer's Order Flow**

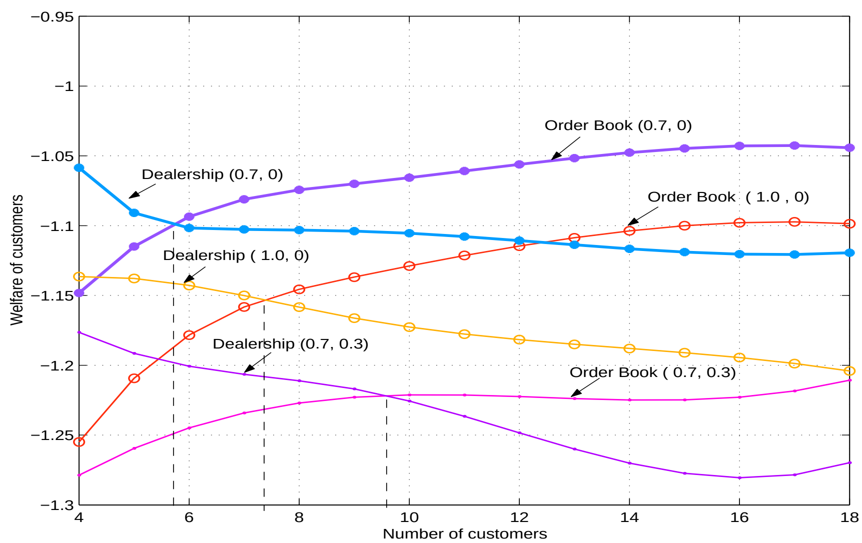

Figure 3: **Smoothed Excess Demand in a Double Auction**

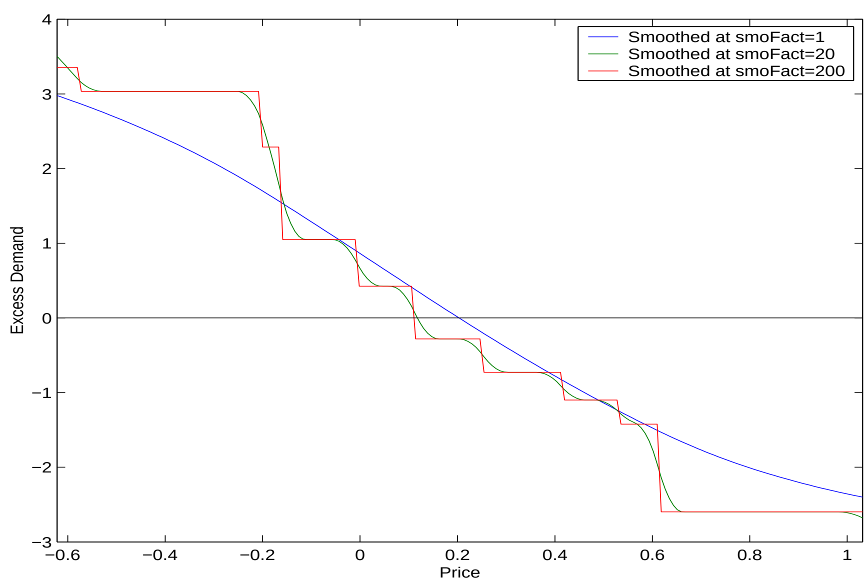

Figure 4: **Customer's Trading Strategy: Bid-Ask**

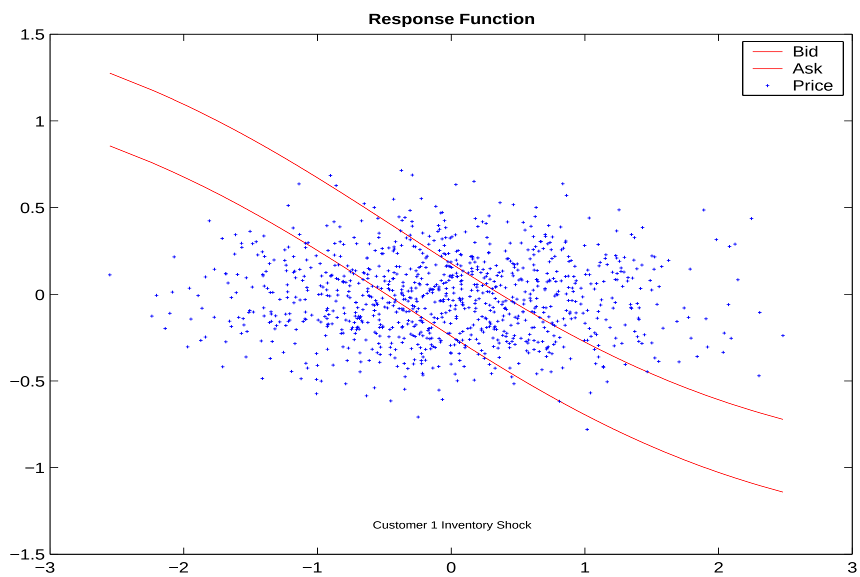

Figure 5: **Customer's Utility**

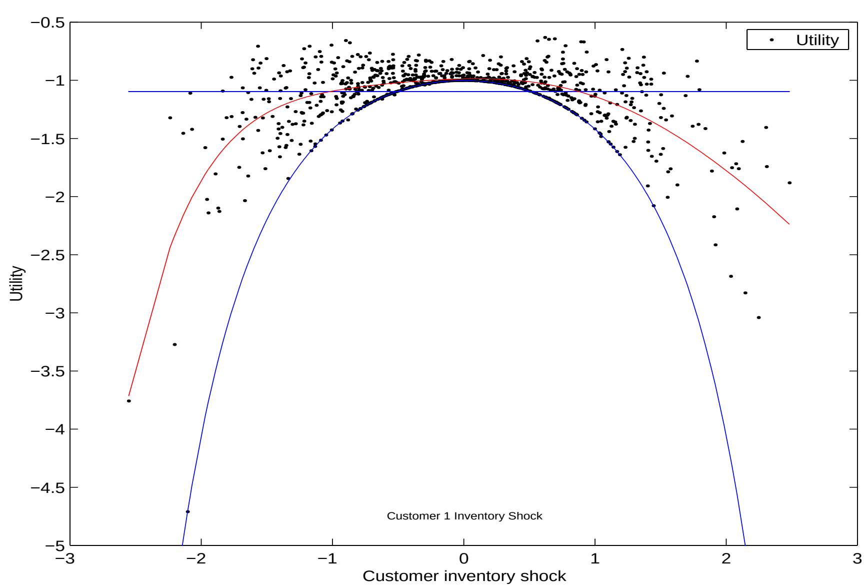

Figure 6: **Risk-Sharing**

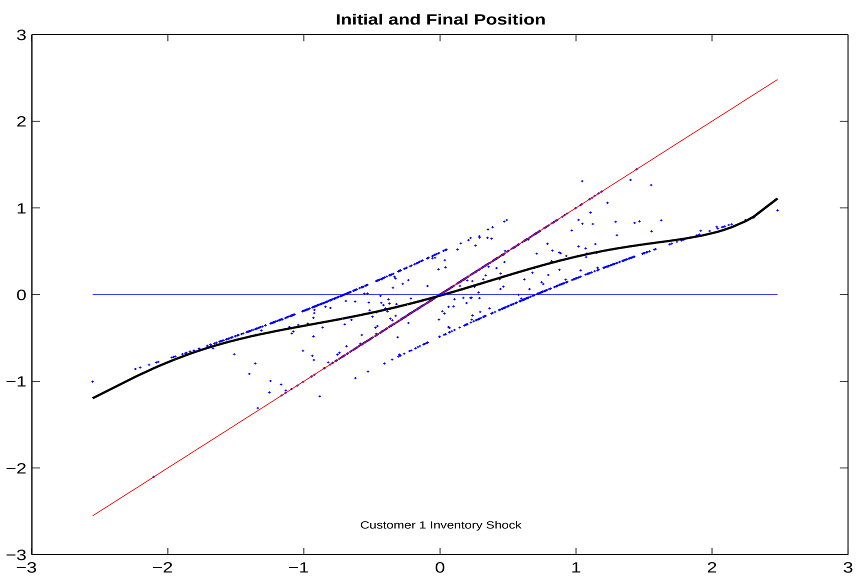

Bank of Canada Working Papers
Documents de travail de la Banque du Canada

Working papers are generally published in the language of the author, with an abstract in both official languages. ***Les documents de travail sont publiés généralement dans la langue utilisée par les auteurs; ils sont cependant précédés d'un résumé bilingue*.**

2002

2002-32	Labour Markets, Liquidity, and Monetary Policy Regimes	D. Andolfatto, S. Hendry, and K. Moran
2002-31	Supply Shocks and Real Exchange Rate Dynamics: Canadian Evidence	C. Gauthier and D. Tessier
2002-30	Inflation Expectations and Learning about Monetary Policy	D. Andolfatto, S. Hendry, and K. Moran
2002-29	Exponentials, Polynomials, and Fourier Series: More Yield Curve Modelling at the Bank of Canada	D.J. Bolder and S. Gusba
2002-28	Filtering for Current Analysis	S. van Norden
2002-27	Habit Formation and the Persistence of Monetary Shocks	H. Bouakez, E. Cardia, and F.J. Ruge-Murcia
2002-26	Nominal Rigidity, Desired Markup Variations, and Real Exchange Rate Persistence	H. Bouakez
2002-25	Nominal Rigidities and Monetary Policy in Canada Since 1981	A. Dib
2002-24	Financial Structure and Economic Growth: A Non-Technical Survey	V. Dolar and C. Meh
2002-23	How to Improve Inflation Targeting at the Bank of Canada	N. Rowe
2002-22	The Usefulness of Consumer Confidence Indexes in the United States	B. Desroches and M-A. Gosselin
2002-21	Entrepreneurial Risk, Credit Constraints, and the Corporate Income Tax: A Quantitative Exploration	C.A. Meh
2002-20	Evaluating the Quarterly Projection Model: A Preliminary Investigation	R. Amano, K. McPhail, H. Pioro, and A. Rennison
2002-19	Estimates of the Sticky-Information Phillips Curve for the United States, Canada, and the United Kingdom	H. Khan and Z. Zhu
2002-18	Estimated DGE Models and Forecasting Accuracy: A Preliminary Investigation with Canadian Data	K. Moran and V. Dolar

Copies and a complete list of working papers are available from:
Pour obtenir des exemplaires et une liste complète des documents de travail, prière de s'adresser à :

Publications Distribution, Bank of Canada
234 Wellington Street, Ottawa, Ontario K1A 0G9
E-mail: publications@bankofcanada.ca
Web site: http://www.bankofcanada.ca

Diffusion des publications, Banque du Canada
234, rue Wellington, Ottawa (Ontario) K1A 0G9
Adresse électronique : publications@banqueducanada.ca
Site Web : http://www.banqueducanada.ca