

Bank of Canada



Banque du Canada

Working Paper 2002-9 / Document de travail 2002-9

**The Microstructure of Multiple-Dealer Equity
and Government Securities Markets:
How They Differ**

by

Toni Gravelle

ISSN 1192-5434

Printed in Canada on recycled paper

Bank of Canada Working Paper 2002-9

April 2002

The Microstructure of Multiple-Dealer Equity and Government Securities Markets: How They Differ

by

Toni Gravelle

Financial Markets Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
tgravelle@bank-banque-canada.ca

This paper is a substantially revised version of a paper included in
Committee for Global Financial Stability Report No. 11 (1999).

The views expressed in this paper are those of the author.
No responsibility for them should be attributed to the Bank of Canada.

Contents

Acknowledgements.....	iv
Abstract/Résumé.....	v
1. Introduction	1
2. Differences Between the Intrinsic Features of Equity and Government Securities.....	2
2.1 Insider information and payoff-relevant private information.....	3
2.2 Term to maturity.....	4
2.3 Inventory-hedging characteristics.....	6
3. Decentralized Markets.....	7
3.1 Decentralized versus centralized markets.....	9
4. Transparency	10
4.1 Centralized markets and transparency	11
4.2 Regulation and transparency.....	12
4.3 Electronic GS trading systems, transparency, and regulation.....	14
5. Some Questions and Future Research Directions	15
5.1 The scale of interdealer and IDB trading	15
5.2 The effects of order-flow dynamics on interdealer trading	17
5.3 Differences in inventory risk-management practices across dealership markets	18
6. Conclusion.....	20
Bibliography	21
Appendix A: Literature Review on the Effects of Transparency in Financial Markets.....	24
Appendix B: The Effects of Greater Post-Trade Transparency in Dealership Markets.....	30

Acknowledgements

I would like to thank Ben Cohen and participants of a BIS working group on market liquidity for their comments. I also would like to thank Walter Engert and Jos Heuvelman for providing invaluable comments. Glen Keenleyside also provided invaluable editorial assistance.

Abstract

Although dealership government and equity securities have, on the surface, similar market structures, the author demonstrates that some subtle differences exist between them that are likely to significantly affect the way market-makers trade, and as such have an impact on the liquidity that they provide. The author reviews some of the concepts recently introduced in the literature, examining multiple-dealer equity markets, and delineates gaps that exist in this literature in terms of its applicability to government securities markets.

JEL classification: G10, G15, G18

Bank classification: Debt management; Financial markets; Market structure and pricing

Résumé

Bien que les marchés de contrepartie relatifs aux actions et aux titres d'État aient en apparence des structures analogues, l'auteur démontre qu'il existe entre eux de subtiles différences, qui sont susceptibles d'influer de façon importante sur la façon dont les teneurs de marché opèrent et donc sur le degré de liquidité du marché. L'auteur passe en revue certains concepts qui sont apparus récemment dans la littérature portant sur les marchés d'actions à courtiers multiples et décrit les lacunes de celle-ci, notamment au chapitre de la modélisation de la structure des marchés de titres d'État.

Classification JEL : G10, G15, G18

Classification de la Banque : Gestion de la dette; Marchés financiers; Structure de marché et fixation des prix

1. Introduction

Until recently, much of the research on the microstructure of securities markets has concentrated on equity markets. It is therefore no surprise that market microstructure research has been heavily weighted towards the development of *asymmetric information models* in which a subset of market participants have private information about the asset's (expected) value. The preponderance of asymmetric information models is itself a result of the bias in market microstructure research towards the price-discovery process in equity markets. The fact that a stock's fundamental value will likely depend on factors that are idiosyncratic to the firm and that changes in those factors are not disseminated to the public on a continuous basis guarantees a dominant role for asymmetric information in the trading process.

Across developed economies, there are two predominant types of equity markets: (i) order-driven, auction-agency markets, and (ii) quote-driven dealership markets. Order-driven markets, or order-book markets, such as the Paris Bourse, are structured as (two-sided) auctions in which there is no intermediary. Incoming orders submitted to the market are either "matched up" with offsetting standing orders previously submitted to the market and placed in an electronic order book, or are themselves placed in the book until an offsetting order is submitted. Although a considerable amount of research has been carried out on both types of equity market structures, a large proportion of the equity *dealership* models has focused on a particular equity market, the New York Stock Exchange (NYSE). Because the NYSE is a specialist (single-dealer) market, most of the dealership models in the market microstructure literature assume a single-dealer equity market rather than a competitive *multiple-dealer* market, such as the Nasdaq in the United States and the London Stock Exchange (LSE). The dearth of microstructure research on multiple-dealer equity markets has been reversed recently by the work of Lyons (1995, 1996, 2001), Saporta (1997), Vogler (1997), Hansch, Naik, and Viswanathan (1998), Reiss and Werner (1998), and Viswanathan and Wang (2000, 2002), to name a few.

An interesting finding of this recent work analyzing multiple-dealer equity and foreign exchange (FX) markets is that the dealers' inventory-management behaviour plays an important role in those markets. This revives a second (earlier) strand of microstructure research known as *inventory models*. In those models, dealers adjust their quote and trading behaviour to restore their inventories to some desired level. The innovations brought forward by this recent line of research are that dealers are assumed to be risk-averse and that dealers trade in a strategic manner to extract information from other dealers.

What is striking in the field of market microstructure research is how little work has been done analyzing the microstructure of bond markets. This has endured despite the size and importance of sovereign government securities (GS) markets, generally the largest segment of a developed country's fixed-income or bond market. Fortunately, GS markets in most developed countries are structured as multiple-dealer markets and function in many ways like multiple-dealer equity markets. There are several important differences, however, between GS and equity dealership markets that may make results garnered from the recent multiple-dealer equity research inapplicable to GS markets. One purpose of this paper is to examine the differences that exist between equity dealership markets and GS markets, to develop a better understanding of the factors that affect the liquidity and efficiency of GS markets. As discussed in Gravelle (1999) and CGFS (2001), well-functioning GS markets play a key role in the maintenance of a stable financial system.

Rather than simply delineate the various aspects and factors that underlie the differences between two securities markets, this paper also reviews some of the recent literature to assess its applicability to GS markets. In doing so, the gaps that exist in the academic microstructure literature are illustrated in terms of modelling the trading structure in GS markets, and some suggestions for future research are offered that would enhance an understanding of the factors underpinning market-makers' incentives to trade.

This paper is organized as follows. Section 2 examines the differences that exist in terms of the intrinsic features of the relevant securities, such as the predisposition to public versus private information, maturity characteristics, and hedgeability. Section 3 examines the subtle structural differences that exist between GS and equity dealership markets. Section 4 focuses on issues related to market transparency. Section 5 describes the factors that affect the level of interdealer trading, such as volatility in customer order arrival and sizes; how this volatility relates to the optimal market structure; non-price competition for order flow; and the interaction of a dealer's market-making activity with their proprietary trading activity. A unifying theme across these issues is that they all impact market-maker behaviour, and as such have a potential impact on the price-discovery process or the level of market liquidity observed in dealership markets. Section 6 concludes.

2. Differences Between the Intrinsic Features of Equity and Government Securities

This section examines some of the differences between equity and government securities that have an effect on their tradeability or on the trading behaviour of market participants (market-makers

and customers alike). The focus is on three intrinsic characteristics: the private information embodied in each type of security, the security's maturity, and the degree of homogeneity within each class of security.

2.1 Insider information and payoff-relevant private information

Although it is quite natural to assume an asymmetric information trading environment for equity markets, this is not a natural assumption in the case of GS markets. It is unclear how much private information exists on the value of GS: prices are dependent on the term structure of the underlying risk-free interest rates. These rates, in turn, depend on macroeconomic factors about which investors do not have private information. In GS markets, private information about the asset's (expected) value—what Cao and Lyons (1998) define as *payoff-relevant private information*—is likely to play a relatively minor role in the agents' trading behaviour. Thus, one of the differences in the intrinsic characteristics of equity and GS is the amount of payoff-relevant private information that each type of security embodies. One can safely assume for modelling purposes that GS hold zero payoff-relevant private information, while equity securities hold some positive amount of private information. This implies that traditional market microstructure asymmetric information models, based on the prevalence of investors who are better informed than dealers, are likely ill-suited to describe the trading environment of the GS market. Moreover, the lack of payoff-relevant information embodied in GS implies that variance in bid-ask spreads (and liquidity) is unlikely to be related to the clustered arrival of informed investors in certain periods.

If payoff-relevant information is of little concern for market-makers in GS markets, then what are the factors that influence their bid-ask spreads and, in turn, their provision of liquidity? Part of the answer seems to lie in a second, older, and, until recently, less-developed branch of the market microstructure literature. In that literature, the focus is on the market-maker's inventory-management behaviour and the assumption is that market-makers stand ready to absorb or, more precisely, *temporally intermediate* temporary imbalances between demand and supply for a security. In these *inventory models*, market-makers charge investors the bid-ask spread in compensation for costs associated with providing immediacy (liquidity).¹ Generally, the bid-ask spread in inventory models depends on the market-maker's level of risk aversion, the asset's riskiness, the market-maker's market power, and, in certain models, their inventory level.²

-
1. A sample of inventory models is given by O'Hara (1995), which includes Ho and Stoll (1983) and O'Hara and Oldfield (1986). Grossman and Miller (1988) model the supply and demand for immediacy in a multiple-dealer setting.
 2. O'Hara and Oldfield (1986) show that the combined uncertainty related to both the "end-of-day" inventory position and the value of that inventory position induces the spread's dependence on the level of inventories. Most inventory models, however, find that the spread is independent of the inventory level, though the placement of the midpoint of the spread is not.

Fleming and Remolona (1999) find that in the market for U.S. Treasury securities, price changes can occur—in tandem with public information releases—without any trading taking place. This is in stark contrast with the asymmetric information models discussed above, where the sequential arrival of information-laden orders moves the price of the stock towards its (new) equilibrium (fundamental) value. Moreover, Fleming and Remolona show that market-makers tend to widen their spreads in response to sharp price movements, a natural tendency for market-makers concerned with the higher inventory risks engendered by greater price volatility. This implies that liquidity in GS markets is closely linked to the market-maker's inventory risk-management costs, which are, in part, subsequently linked to variations in the market-maker's level of risk aversion and GS price volatility.³

2.2 Term to maturity

Section 2.1 argued that one of the intrinsic differences between equity and government securities was the amount of payoff-relevant private information that each embodied. Another intrinsic difference between these two types of securities, which has an impact on the trading decision of market participants and, in turn, on relative market-liquidity characteristics, is term to maturity. Stocks have an infinite maturity while government debt securities have a finite maturity. The GS fixed-maturity date generates two classes of investors for the security. The existence of two investor types, in turn, has an impact on the number of securities available for trade, which has an impact on the security's liquidity characteristics. The existence of two types of GS investors implies that there is a *floating supply* of the security available for trading before its maturity date that is less than the total amount issued to the public.

A finite maturity structure for a security implies that investors have the option of liquidating their position in that security at a *known* maturity date, in addition to liquidating their position sometime before that date via the secondary market. This option is not available for equity investors, who generally must liquidate the stock in the secondary market and, in turn, incur some market trading costs. Moreover, the equity investor's holding period, as is normally the case for all utility-maximizing agents, is conditional on the stock's expected return at any given time, which in turn depends on the stock's price and the agent's private information and expectations at that time.

3. Fleming and Remolona (1999) also show that, in the U.S. GS market, public information precipitates above-average trading volumes theorized to be associated with uninformed liquidity traders (i.e., traders who trade in reaction to public information to rebalance their portfolio, for example, rather than trade on private information).

On the other hand, a government debt security provides the investor with the option of fixing their holding period to a known date. This temporal optionality drives a wedge in the pool of available GS investors, thus generating two types of investor: buy-and-hold and trading market participants. The predetermined liquidation date and liquidation value of the debt security gives GS investors the option of *locking away* those securities in their portfolios until maturity, which reduces the supply of the security available for trading purposes. Specifically, a large pool of investors not only derives utility from a fixed-income instrument's (expected) return but also from the fact that the instrument's cash-flow stream and liquidation value are *known* (ex ante).⁴ In sum, the fixed maturity date of a GS creates the possibility that the supply of the GS available for trading—the *floating* or *effective supply*—is less than the total (outstanding) supply, which is not the case for equities.⁵ Moreover, since the value of the GS's predetermined holding period increases as its maturity date approaches, the amount of the security in the hands of the buy-and-hold investors will tend to increase over this period, causing the floating supply of the GS to fall over time.

A decrease in the floating supply has a detrimental impact on the liquidity of the security both in terms of trading intensity and bid-ask spreads. A smaller number of securities in the hands of trading market participants tends to have a direct negative impact on the trading intensity of the security.⁶ In turn, market-makers find it more difficult to carry out their inventory-control activities as trading activity decreases with the drop in floating supply. Specifically, as floating supply decreases, the implicit costs incurred by market-makers increase, since they must wait longer for the arrival of rebalancing orders that move them back towards their desired inventory level or increase their search efforts for such orders.⁷ In other words, dealers find it increasingly costly to temporally mediate transactions.

-
4. Singleton (1996) describes how institutional investors that revalue their portfolios on a regular basis tend to be active market participants, while those that do not tend to be buy-and-hold investors. The latter may be discouraged from selling the bond (because of accounting practices) at a price that does not lie between the purchase and par (maturity) value and, as such, tend to be buy-and-hold investors. Singleton argues that a greater proportion of buy-and-hold investors reduces the floating supply of bonds in the market.
 5. Some firms may choose to hold (or repurchase) a proportion of their own stock, thus allowing the amount of the stock available for trading to differ from the total amount of issued stock. Since the firm's objective in holding some of its stock is to maximize its price, however, it avoids reducing the supply of the stock to the point of affecting liquidity. Further, it should be clear that a floating supply of GS arises from the utility-maximizing behaviour of the investors, and not the issuer.
 6. Amihud and Mendelson (1986) show that there is in equilibrium a positive correlation between trading intensity (defined as the number of transactions observed over a certain period) and trading volume.
 7. Any permanent decrease in the security's liquidity will increase the costs associated with actively trading the instrument and therefore increase the likelihood that investors holding the security will hold it until maturity. This drop in liquidity tends to decrease the floating supply of the security, which, in turn, tends to reinforce the decrease in liquidity. Moreover, because of positive participation externality effects on other traders (noted by Harris 1993, among others), any decrease in the floating supply may have a greater than one-for-one negative impact on the customer activity observed by the market-makers.

2.3 Inventory-hedging characteristics

One of the differences between GS inventories and equity inventories is that there are a far greater number of instruments available to hedge GS inventories. The market-maker's inventory risks associated with holding a specific government bond can be hedged using a variety of instruments: they can borrow or lend the same security in the repo or lending markets, they can hold the opposite position of a nearly identical bond (one with a similar duration), or they can offset this position with a position in the related futures contract. Alternatively, they could simply try to rebalance their inventory position to reduce their risks.

For equity market-makers, avenues for hedging inventory balances in a specific stock are, for all but the most actively traded securities, much more limited. Often, the only inventory risk adjustment available to them is to rebalance inventory. They cannot, in general, hedge their inventory by taking an opposite position in a similar stock, nor can they hedge using a futures contract, because there are no instruments (such as another stock or a futures contract) sufficiently correlated to the stock's price. This contrasts with GS, where the yield movements for most cash instruments move in a correlated fashion. For example, when macroeconomic news is released, most of the GS yields (within a certain sector of the yield curve) move in the same direction and by a nearly equivalent amount (in yield terms), such that the (historical) relation between those securities' yields remains relatively constant. With this high correlation in yield movements, market-makers can easily find a large list of other GS in which to take an opposite position to offset an inventory imbalance in a specific security. This homogeneity in yield movements can even be applied to the greater set of cash fixed-income instruments. That is, in general, even bonds issued by a variety of non-government issuers will move in a manner that holds their historical yield spreads relative to government bonds fairly constant (at least at a daily frequency). This allows market-makers in a specific corporate bond, for example, to sell short a GS as a hedge when they are holding a long position in a corporate bond. Moreover, those same market-makers can often use a position in a GS futures contract to hedge their position in the corporate bond.

In sum, the greater homogeneity among GS in terms of yield dynamics implies a greater ease of inventory price hedging for the GS market-makers relative to their counterparts in equity markets. This enables GS dealers to endure greater extremes in their inventory balances than equity market-makers. This point is discussed further in section 5.

3. Decentralized Markets

As opposed to single-dealer markets, such as the NYSE, market-makers in multiple-dealer markets must directly compete for their share of the order flow.⁸ Moreover, competing dealers have the option of laying off unwanted inventory positions by trading with other dealers, which enables them to share their inventory risks across the market-maker community, rather than be constrained to rebalance their inventories by awaiting the arrival of public investor orders, as is generally the case for the NYSE specialist.

As a result, at least three facets of multiple-dealer markets cannot be adequately captured in specialist-based theoretical models. First is the competition for customer order flow that exists between dealers. Second is the existence, in most multiple-dealer markets, of two parallel trading environments: a public trading environment, where customers trade exclusively with market-makers, which we call the public sphere, and an interdealer trading environment, which we call the interdealer sphere. This implies that a dealer in multiple-dealer markets has the option of managing their inventory via transactions with other dealers, rather than waiting for the arrival of offsetting customer orders. The third facet is the further segmentation of the interdealer trading environment into two trading mechanisms. In this environment, the dealers can choose to either trade bilaterally with each other or trade indirectly with each other anonymously via an interdealer broker (IDB) system, which resembles the auction-agency trading structure of many order-driven exchanges.⁹ Although these three facets of multiple-dealer markets are generally applicable to both equity and GS dealership markets (and FX markets as well), in certain equity or GS markets the services provided by an IDB system are offered by an electronic auction-agency system, instead.¹⁰

Given the relatively complex nature of the two-sphere trading environment for multiple-dealer markets (also called a two-stage trading environment), it is perhaps not surprising that research on multiple-dealer markets is undeveloped. A sample of multiple-dealer models includes Ho and Stoll (1983), Leach and Madhavan (1993), Perraudin and Vitale (1996), Vogler (1997), and Naik, Neuberger, and Viswanathan (1999). Various problems with these models, however, limit their

-
8. Specialists face competition of a similar nature. They must compete with the order book, which lists all standing limit orders. However, the competition is on, to some extent, the specialist's terms. That is, since the specialist observes all the order flow for the stock (including that flowing to the order book), they can set their bid and ask price based on this knowledge.
 9. See O'Hara (1995) for a review of order-driven auction-agency markets.
 10. Instinet provides in essence the same anonymous interdealer trading services as an IDB system for the Nasdaq equity market. The EuroMTS system in European GS markets provides anonymous interdealer bond trading (CGFS 2001).

use in capturing the various complex features of multiple-dealership models. For example, even though Ho and Stoll (1983) were the first to model interdealer trading, they assumed a high degree of transparency in the trading environment. Specifically, they assumed that transactions, and thus dealer inventories, are public information. They also assumed that both customer trades and interdealer trades are carried out in an identical manner via the public sphere, which, as described above, does not capture the essence of the parallel trading environment that exists in most multiple-dealer markets. Although the more recent theoretical work on dealership markets has relaxed some of these simplifying restrictions, it has in general failed to combine both private information and risk-averse market-makers in one model, features essential to GS markets, as discussed above. Moreover, most models that do examine private information focus on payoff-relevant private information, which is not necessarily a feature of GS markets. Section 3.1 considers how the decentralized nature of GS markets differentiates them from equity dealership models, and how the opaque nature of decentralized markets introduces a role for payoff-irrelevant information in the strategic behaviour of market-makers.

As stated in section 2, in contrast to equity securities, payoff-relevant private information about the value of a GS is unlikely to be important in the provision of market liquidity and the price-discovery process. The multiple-dealer market structure described above does, however, allow information asymmetries to be a prevalent feature of GS markets. That is, although GS market participants are unlikely to have superior (or private) information about a security's payoff (or fundamental value), that fact does not preclude certain agents, namely market-makers, from having private information about the state of the trading environment, such as customer order flow, that will help them to better predict the intervening price movements.

Cao and Lyons (1998) call this class of information *payoff-irrelevant private information* and show that it, coupled with a market-maker's risk aversion, is an essential determinant of the price-discovery process in decentralized multiple-dealer markets, such as GS markets.¹¹ Payoff-irrelevant private information is, in principle, also prevalent in equity dealership markets. But the impact of this type of information asymmetry on the trading process is likely to be less than in GS markets, given public investor access to payoff-relevant private information and, as described below, the equity markets' greater degree of order-flow transparency. Using market microstructure terminology, this implies that market-makers are much more likely to be the "informed agent" in GS dealership markets than their counterparts in equity dealership markets.

11. Lyons (2001, chapter 4) provides an excellent discussion of the drawbacks of the two traditional approaches to market microstructure theory: information and inventory models. He also shows how models, based on payoff-irrelevant information-induced asymmetries and dealer risk aversion, bridge the gap between the two traditional approaches.

3.1 Decentralized versus centralized markets

Multiple-dealership equity markets (such as the LSE and Nasdaq) are centralized, whereas multiple-dealership GS markets are decentralized. Decentralized and centralized markets are distinguished by the amount of information that is available to the public (both customers and dealers alike) on a consolidated basis. For example, even though Nasdaq operates as an over-the-counter (OTC) market, it is nonetheless linked together electronically, so that price and trade information can be viewed on a consolidated basis by market participants. This includes pre-trade information (bid and ask quotes across dealers) and post-trade information (data on completed trades in the market).

Here, “consolidated basis” means that price and trade data from the spectrum of dispersed dealers are available on a single screen. For example, multiple-dealer quotes on Nasdaq are generally available to the public in a consolidated format, such as a single Bloomberg page.¹² This is not generally the case for GS markets. GS markets (and fixed-income markets more generally) are also multiple-dealer markets, but in contrast with the Nasdaq and the LSE they tend not to be linked electronically. As a result, and in contrast with multiple-dealer equity markets, investors in GS markets cannot easily ascertain the best, most current bid-ask spread being offered by dealers. In principle, the only way for an investor to ascertain which dealer has the best quote in a decentralized market would be for them to contact each dealer directly. Therefore, given the decentralized nature of GS markets, it is possible for simultaneous transactions to occur at different prices and, more importantly, at prices other than the best available price across a spectrum of market-makers.

In the U.S. and European GS markets, electronic trading systems such as TradeWeb and BondVision have recently appeared and have, in effect, increased the degree of centralization in those markets, since they allow investors to view multiple-dealer quotes on one screen (i.e., in a consolidated format).¹³ These electronic trading platforms allow investors to solicit quotes from a number of participating dealers *simultaneously* (without the traditional need to contact them *sequentially* over the telephone) and to trade electronically with the dealer of their choice.

-
12. Nasdaq offers three levels of information access. Public investors, via their brokers or information vendors, have access to level 1 terminals, which allows them to observe (real-time) best bid-ask prices. Level 2 screens, which are available to non-dealer brokers, allow them to see multiple market-maker bid and ask quotes. Nasdaq market-makers have access to level 3 terminals (screens), which enables them to enter their bid and ask quotes and the volume of the security to sell at those quotes.
 13. These systems operate in the public sphere of GS markets and should not be confused with electronic fixed-income interdealer broker systems, such as EuroMTS or ESpeed, which cater to interdealer trading.

Although these electronic multiple-dealer systems offer significantly improved time-savings and operational efficiency (particularly in terms of the back-office efficiencies offered by straight-through processing), these quote-driven systems do not alter the way in which investors interact with dealers. In essence, these systems are simply an automation of the telephone-based investor-to-dealer interaction. Moreover, GS investors (and dealers) participating in these electronic trading platforms remain concerned with minimizing the price-impact (or trading) costs of large GS transactions. As section 4 will show, this degree of centralization and investor concerns about trading costs play an important role in determining the appropriate level of transparency in dealership GS markets.

4. Transparency

Generally, a market increases its transparency when more data on its internal trading processes are publicly available. As with market liquidity, however, measuring the degree to which a market is transparent is not easy. Market transparency is multi-dimensional and goes beyond simply measuring quote and trade data. Although a market participant's trading behaviour depends in general on the quotes they observe and on the execution price of the most recent transaction(s), their behaviour also depends on trader identity, trade size, how long it took to execute the trade, the size of any posted limit orders, etc.¹⁴ Thus, increases in publicly available data on dealer quotes or transaction prices would imply an increase in transparency only if the other dimensions of market transparency were held constant across the regimes.

Existing research on market transparency tends to be arranged around questions relating to *pre-trade transparency* and *post-trade transparency*. In the case of multiple-dealer markets, it is helpful to make the following distinctions. A market is classified as being pre-trade transparent when traders (customers and dealers alike) can directly view *all*, or at a minimum the best, firm bid-ask quotes offered by the spectrum of market-makers. Therefore, a key feature of a pre-trade transparent market is that market participants can, at all times, see all prices that are available and have access to, or can trade at, the best available price. In some of the literature, a distinction is made between the public display of the best bid-ask quote versus the display of all quotes, with markets that provide consolidated access to all quotes being described as more transparent than those that simply display the best prices available. Post-trade transparent markets, on the other hand, are classified as being those that report *all* completed trades to the public immediately.¹⁵

14. See Scalia and Vacca (1999) for an example of research into the effects on market liquidity of a transparency regime change caused by curtailed information on trader identity.

15. More generally, pre-trade transparency refers to the amount of information provided about available bid and ask prices and volumes, and how widely accessible this information is to all market participants. Post-trade transparency refers to the extent of information available on the completed transactions (including price, volume, trader identity), and how fast this information is released to market participants. See Pagano and Röell (1996) for some post- and pre-trade transparency definitions.

4.1 Centralized markets and transparency

In terms of pre- and post-trade transparency, equity dealership markets such as Nasdaq and the LSE are superior to GS markets. This stems in part from the centralized nature of the equity dealership markets, where dealer quotes are reported on a consolidated basis. The existence of a consolidating (electronic) display system also allows for greater disclosure of marketplace information, since those systems make it more *feasible* to collect and in turn transmit or disseminate pre- and post-trade data. Given the decentralized nature of most GS markets, consolidated pre-trade quote information is generally not available to public market participants. Therefore, key factors underlying the degree of transparency in dealership markets are the existence of a system that consolidates and collects (links electronically) the individual dealer quotes, and the accessibility of those data by all potential market participants.

As noted in section 3, the recent arrival of electronic multiple-dealer systems in the public sphere of GS markets has increased the degree of consolidation that exists in U.S. and European securities markets. One feature of those systems, which enhances the level of pre-trade transparency in the GS markets, is their ability to transmit the indicative quotes posted by dealers participating in the system. Consequently, for those investors who have access to those systems, there is a higher degree of pre-trade transparency than is generally available in other GS markets.

In most GS markets, IDBs provide a degree of centralization that, when viewed from the public sphere, are completely decentralized. That is, the decentralized nature of GS markets causes the services provided by the IDBs in those markets to take on added importance when compared to IDBs that serve dealers in centralized equity markets. Under the assumption that dealers compete for a share of market-wide order flow, the dealers will tend to minimize the difference between the (transparent to all dealers) IDB quotes and those offered to investors. This, ultimately, should reduce the probability that simultaneous public orders are executed at significantly different prices from each other, thereby improving the degree of centralization in GS markets.

The fact that market-makers in GS markets play a much greater role in the price-discovery process than their counterparts in dealership equity markets (such as the LSE and Nasdaq), which tend to be restricted to providers of immediacy, means that GS market-makers rely on interdealer trading more for information purposes than for risk-sharing services, as is generally the case for equity market-makers. Therefore, because of the relatively infrequent customer orders, GS dealers not only temporally intermediate trades and supply liquidity, but also, through their trading on IDBs, enhance the price-discovery process in GS markets. This implies that the efficiency of the

GS market is underpinned by the ability of the GS market-makers to ensure an effective price-discovery process and to minimize the transactions costs arising from market impact.

Although the existence of a consolidating electronic system in equity dealership markets such as Nasdaq and the LSE makes it easier to disclose post-trade information to the public, ultimately the regulatory framework, given the market-maker incentives described in section 4.2, ensures that this information is disclosed to all market participants in a timely fashion. For example, the Securities Exchange Commission (SEC) mandates the quick reporting of all equity trades in the United States. For Nasdaq, all trades are disclosed immediately, no later than 90 seconds after completion. Other jurisdictions, however, do not hold the same views as the SEC about the appropriate level of post-trade transparency in equity dealership markets. The LSE has a lower level of post-trade transparency, in which small trades and all trades for FTSE-100 stocks are reported immediately, while the reporting of larger trades can be delayed by up to 60 minutes after completion.¹⁶ The different approach taken across equity markets highlights the long-standing debate about the appropriate level of transparency in securities markets.

4.2 Regulation and transparency

The main goal of financial market regulators is to improve investor welfare. As such, equity market policy-makers have in general sought to uniformly increase the level of transparency in the belief that it will protect investors from being unfairly treated during the trading process, and in the belief that transparency unambiguously improves the informational efficiency and functioning of markets. However, there is increasing evidence that a trade-off exists between greater transparency and market quality as measured by market liquidity and trading costs. (Appendix A reviews the academic literature that deals with transparency, focusing on transparency in dealership markets.) As outlined in Appendix A, as well as in Madhavan (2000) and Ganley et al. (1998), the evidence across many studies indicates that greater transparency seems to have a “Laffer curve” effect on markets. That is, although there is generally a consensus that some positive amount of quote and trade disclosure improves market efficiency and liquidity, too much transparency will actually negatively impact the quality of a market. Moreover, there is increasing evidence in the literature that a “one size fits all” approach to transparency regulation is not appropriate when considering the amount or extent of transparency that is optimal for different securities or different market structures.

16. See Ganley et al. (1998) for more details on various LSE transparency regimes. Note that, since October 1997, the LSE has become a hybrid market in which the FTSE-100 stock now trades on an order-book system. (One can still trade these shares via a market-maker, but now those transactions are called “upstairs” trades.)

There is a particular tension between post-trade transparency and market liquidity and trading costs. As discussed in more detail in Appendix B, greater post-trade information disclosure in dealership markets may not benefit the overall market, since it impinges on a market-maker's ability to carry out its inventory risk-sharing activity. There is a natural incentive for market-makers to seek to delay the disclosure of completed trades to the broader market.¹⁷ These market-maker incentives help to explain the different approach to post-trade equity market transparency taken across jurisdictions. That is, they help to explain the arguments put forward in defence of delayed trade reporting for large orders on the LSE, as discussed in O'Hara (1995, 258–59) and Board and Sutcliffe (1995).¹⁸

What is the optimal level of post-trade transparency in equity and GS dealership markets? The differential disclosure treatment offered for large trades on the LSE offers some guidance. Large orders tend to be particularly sensitive to so-called “Hirshleifer revelation risks” (see Lyons 1996 and Naik, Neuberger, and Viswanathan 1999). That is, increases in post-trade transparency will leave market-makers with less time to unwind their inventory positions before the broader markets become aware of the trade. As market-makers find it increasingly costly to manage their inventory because of increased revelation risks, they must pass on the additional costs to investors. As Appendix B indicates, the larger the trade and/or the less frequent the trading activity, the more susceptible market-makers are to *revelation risks*. In the end, investors who must transact relatively large orders are made worse off as they must trade in a market that charges a higher price (in terms of the bid-ask spread, for example) to access the market liquidity services provided by market-makers.

Public investors in GS markets are predominantly large institutional investors, whereas there is a greater preponderance of small retail investors in equity markets. As a result, it is likely that less

-
17. Regulation is sometimes also necessary to ensure some level of pre-trade transparency. However, market-makers often have an incentive to disclose pre-trade information to the broader market, to increase market share or to attract offsetting investor orders that enable the market-maker to better manage their inventory positions (see Lyons 1996 on the latter issue). Consequently, the market-makers' incentives are such that little regulatory prodding is necessary to ensure disclosure of pre-trade information.
 18. Flood et al. (1997) state that the enforceability of a certain level of post-trade transparency is directly linked to the degree of centralized order processing. They argue that electronic or floor-based trading, such as the NYSE or Paris Bourse, provides a higher level of post-trade transparency than the LSE or Nasdaq, because the order flow must pass through a centralized exchange. They indicate that it is difficult to enforce strict post-trade disclosure in dealership markets because larger trades tend to be conducted bilaterally over the phone, and must be reported by participants to the authorities. Although the findings of abundant protected trading on the LSE by Board and Sutcliffe (1995) and Reiss and Werner (1998) is expected, given the LSE's delayed disclosure rules, the findings of Porter and Weaver (1998), which indicate that dealers on the Nasdaq systematically delay trade reporting, is surprising and inconsistent with the reporting guidelines of the SEC.

post-trade transparency in the public sphere of GS markets than is currently observed in equity dealership markets would be appropriate.

4.3 Electronic GS trading systems, transparency, and regulation

Electronic debt-trading systems have been developing business plans to enter the public sphere of the Canadian fixed-income market. As described in section 3, these electronic trading systems represent an automation of the current phone-based method of trading in the public sphere of the Canadian GS market.

Additionally, Canadian securities regulators (the Canadian Securities Administrators) recently published a document that proposes a regulatory framework to guide the operation of electronic securities trading systems in Canada. The goal is to allow new electronic systems to compete with the incumbent trading venues and yet minimize market fragmentation. Included in the regulatory framework are rules that set out the level of post-trade transparency that the electronic systems must adhere to. In the initial form of the proposal, the application of the transparency rules did not differ materially across equity and debt market electronic trading systems. Both equity and GS market electronic systems were mandated to disseminate real-time post-trade information market-wide. As already noted, however, GS markets differ in many ways from equity dealership markets and are likely to require different regulatory treatment. This is particularly true for transparency rules, where the quality of GS markets may be negatively affected by too much post-trade information disclosure. Thus, even though the architecture of certain equity and GS electronic systems may be identical, this does not mean that the regulatory framework guiding their operations should be identical. (Note that the final form of the regulatory framework for electronic trading systems in Canada, which took effect December 2001, does have different transparency requirements across equity and GS electronic systems.)

Allen, Hawkins, and Sato (2001) make a similar point. They explain that it is increasingly important for policy-makers to recognize that greater transparency has a different impact on market quality across different markets. Specifically, they note that although the understanding about market transparency is incomplete at present, it is known that there is not a “unidirectional relationship between transparency and quality of markets.”

In summary, although the advent of electronic trading systems increases the *feasibility* of highly transparent markets, policy-makers should nonetheless consider different transparency arrangements that vary across market structures, market segments, and trading environments, and consider that certain market participants benefit often at the expense of other participants when transparency is increased.

5. Some Questions and Future Research Directions

This section considers various issues that have not received much attention in the microstructure literature, but which are important to increase our understanding of the factors that affect market liquidity and price discovery in dealership markets.

5.1 The scale of interdealer and IDB trading

One aspect of dealership markets that has received little attention in the literature, at least until recently, is the existence of interdealer markets that operate in parallel with the customer or public market. In particular, there has been little research to examine the microstructure of interdealer markets in which dealers have a choice of trading bilaterally with other market-makers or trading in an IDB market that resembles a screen-based “order-driven” auction-agency market.

Recent work by Viswanathan and Wang (2000), however, shows that, under certain market structures, dealership markets *with* interdealer trading are preferred by customers. Viswanathan and Wang (2002) explain that, when order flow can be separated into the small and large, customers prefer to trade in hybrid markets, where small orders are routed to an order-driven market and large orders to a dealership (or “upstairs”) market.

A broader question, of importance to asset markets more generally, concerns the scale of interdealer and IDB trading across different dealership markets. Table 1 illustrates the scale of interdealer trading across various GS markets for the FX market and the LSE. It shows that although the proportion of interdealer trading (relative to total trading volumes) carried out across markets does have some consistency, except for the dollar/mark FX market, there is some variation in the usage of IDBs by dealers across different dealership markets. Why do dealers in markets such as those for the GS tend to rely almost exclusively on IDBs when trading among themselves, while dealers in the FX markets or the LSE trade bilaterally to a greater extent? Lyons (1996) suggests that IDBs play a role in the information-extraction process for FX dealers. Lyons posits that dealers endeavour to ascertain market-wide order-flow information by observing the order flow passing through all IDB systems. In a decentralized market this information is vital, since it signals information that market-makers have gathered from public orders or from other private sources.¹⁹

19. Vitale (1998) confirms this hypothesis for the U.K. gilt market. He shows that the total quantity of interdealer trading has a significant positive impact on the individual market-maker’s transacted price.

Table 1: Government Securities Market Microstructure Characteristics

Country	% total trading volume that is inter-dealer trading	% of total interdealer trading via interdealer brokers
U.S. Treasury (1997)	50	95 ^a
Government of Canada (2001)	47	83 ^b
Japan (1997) ^c	60	60
U.K. (1997)	47	98 ^d
Other security types		
FX ^e	80	50
LSE ^f	40	40–60

- a. Fleming and Remolona (1999) indicate that over 90 per cent of interdealer trading is via IDB, while Dattels (1995) indicates that IDB trading is near 99 per cent in the U.S. Treasury market.
- b. IDB trading has grown in Canada over the years. For example, in 1991, IDB trades accounted for 50 per cent of interdealer trading. Even in 1997, IDB trading accounted for 75 per cent.
- c. Approximation based on total fixed-income trading, of which GS trading accounts for 95 per cent. Source: Miyanoya, Inoue, and Higo (1997).
- d. Dattels (1995) indicates that IDB trading accounts for 96 per cent of interdealer trading, while a recent BIS-CGFS survey indicates that it accounts for 100 per cent.
- e. U.S. dollar/German mark FX trading. Source: Lyons (2001).
- f. It is more precise to give the range for the share of interdealer trading as being 24 per cent to 53 per cent. Hansch, Naik, and Viswanathan (1998) indicate that it is 53 per cent and Reiss and Werner (1998) state that the LSE estimates interdealer trading to be 30–50 per cent of total LSE trading volume, while they themselves estimate that it is in fact 25–27 per cent for their sample period. This discrepancy occurs because Reiss and Werner control for double counting. They also provide data on the proportion of IDB trading in the LSE interdealer market. Vogler (1997) notes that, based on 1994 data, interdealer trading is 40 per cent of total trading.

This analysis does not explain why dealers choose to use IDBs to *execute* interdealer trades instead of bilateral interdealer transactions. As well, it still must be determined why proportionally more IDB trading in GS markets is observed than in equity dealership markets.²⁰ Does the scale of interdealing trading or IDB trading have an effect on the liquidity offered by

20. The author is aware of only one example of research investigating a dealer's motivation in choosing an interdealer trading mechanism: Saporta (1997). Many of the issues discussed in section 5 are examined by Saporta using a theoretical three-stage model. Although Saporta does examine the relationship between asset-price volatility and the dealer's choice of interdealer trading venue, she does not examine the relationship between interdealer trading and public order-flow dynamics.

dealers to customers (in terms of bid-ask spread, depth, or volume)? How does a decentralized multiple-dealer market arrive at a particular level of IDB usage? These questions are integral to the liquidity of most GS markets, since interdealer trading has an impact on a dealer's ability to provide liquidity to customers.

5.2 The effects of order-flow dynamics on interdealer trading

A better understanding is required of the effect of order-flow dynamics on interdealer trading in dealership markets. Recent empirical work by Hansch, Naik, and Viswanathan (1998) presents empirical evidence that sheds some light on this for market-makers on the LSE. They show that interdealer trading facilitates the dealer's inventory risk-management process. This allows dealers to take on large inventory positions that they would otherwise be unwilling to take if they could rebalance positions only through the arrival of offsetting public orders. Specifically, they show that dealers are more likely to engage in interdealer trading when their inventories are high or when they hold extreme, relative inventories.²¹

It follows that an important factor affecting the proportion of interdealer and, to a greater extent, IDB trading in dealership markets is the arrival rate of customer orders and the variation in size and direction of those orders. Risk-averse dealers, subject to greater order-arrival variability and/or order-size variability, will likely require a greater amount of inventory risk-management services. As Ho and Stoll (1983) show, interdealer trading will occur as dealers choose between the uncertainty of public trade arrival and the certainty of interdealer trading. That is, when customer-to-dealer trading intensity is low, risk-averse dealers are not likely to witness a rebalancing order from the public sphere, and thus will tend to rebalance in the interdealer market.

Moreover, dealers are more likely to require the anonymous trading offered by IDBs (rather than direct "name give-up" interdealer trading) if the order's size is large enough to move markets. One possible explanation for the greater use of IDBs in GS securities markets than in other markets (see Table 1) is that GS market-makers receive orders that are, relative to their desired (risk-adjusted) inventory level, on average larger as well as less frequent than their counterparts in dealership equity markets. Both Proudman (1995) and Vitale (1998) note that the U.K. gilt market is characterized by large and infrequent order flow relative to that observed in the U.K. equity market.

Another possible explanation is that public orders are more likely to arrive in bunches for GS, since trading volumes jump immediately after macroeconomic news releases. Although there are

21. Reiss and Werner (1998) report similar results.

news events that may cause similar jumps in trading volumes for certain stocks, most firms now release material information after the close of the exchange trading hours, allowing for a more balanced order flow at the opening of the following trading day. This implies that GS market-makers tend to have a greater need for inventory rebalancing given their greater susceptibility to news-driven, one-way public order flow (which is driven by uninformed, public liquidity traders). Equity market dealers also engage in preferencing arrangements (or have vertically integrated operations), which reduces the variance in their (public) order arrival rates.

The difference in order-flow dynamics observed by dealers in dealership equity markets versus those in GS markets may be one explanation for the important difference between the interdealer trading structures of these markets. What is fundamentally driving the difference in order-flow dynamics between GS and equity markets? Institutional investors form the majority of the customer base in GS markets, unlike in equity dealership markets, where individual investors participate actively. Institutional investors, such as pension funds, mutual funds, and insurance companies, are more sophisticated, which allows them to better infer the effect of macroeconomic news on GS prices. This gives them an advantage over individual investors when negotiating and/or trading with sophisticated market-makers (who, in general, have a dedicated research wing investigating the relationship between macroeconomic variables and bond prices). Given the fact that these institutional investors pool individual investor demands for GS, their transactions will, on average, be large and, since this type of customer dominates the GS sphere of customers, the order flow observed by GS market-makers will be relatively large, on average, and less frequent. Moreover, fund managers are evaluated according to common benchmarks. This implies that the trading behaviour of institutional investors is guided by movements of assets relative to the common benchmarks' indices. This, too, will tend to cause their trading activity to be correlated (i.e., herd-like), particularly if there are strong movements in the benchmark indices.

5.3 Differences in inventory risk-management practices across dealership markets

In the previous section it was argued that interdealer trading activity might be driven in part by the public order-flow dynamics observed by dealers, and that differences in interdealer trading activity and structure across dealership equity and GS markets might be the result of differences in public order-flow characteristics in those markets. This section discusses why differences in the way market-makers manage their inventory positions are also likely to exist across GS and equity markets.

The ability of market-makers to take on large public orders (thus offering liquidity) depends in part on the way they manage their inventory positions. As discussed earlier, the interdealer market provides an outlet for dealers to share inventory risks with other dealers. However, the intensity with which dealers manage their inventory imbalances varies across markets. Hansch, Naik, and Viswanathan (1998) note that mean reversion in the LSE market-makers' inventories is stronger than mean reversion in the NYSE specialists' inventories (2.5 days versus 7.3 days), while Lyons (1995) indicates that a large U.S. dollar/mark FX dealer's mean reversion is even greater (with an inventory half-life of 10 minutes). Anecdotal evidence gathered from the Government of Canada (GoC) securities market seems to indicate that mean reversion is much weaker in GS markets than in equity and FX markets. In particular, GS market-makers do not seem to manage inventories as intensely as FX market-makers. In contrast, GS market-makers tend to hold significant inventories over weeks and the desired level of those inventories seems to change over time.²²

Related to this issue is the observation that risk-averse GS market-makers need not adjust their quoted spreads (reflecting inventory-control effects) as aggressively as equity market-makers, since their inventory can be hedged against price movements (using futures contracts or a mix of offsetting GS).²³ Some evidence in support of this observation is found in Vitale (1998). In contrast with the results in Hansch, Naik, and Viswanathan (1998), showing inventory effects on LSE market-makers' transaction prices, Vitale shows that imbalances in inventories do not influence U.K. gilt market-maker transaction prices, and suggests that this is the result of the availability of several hedging avenues for GS market-makers.²⁴

Taken together, these observations seem to imply that GS market-makers can endure a greater divergence of inventories (or more extreme inventory levels) and, given the lower inventory rebalancing needs, are less likely to engage in interdealer trading. This contradicts the data in Table 1 that show that GS market-makers rely more on interdealer trading than do LSE market-makers. We conjecture that, for GS market-makers to competitively supply liquidity to the public sphere, they must be able to *both* endure greater inventory imbalances *and* engage in interdealer trading. To be more precise, even though GS market-makers require more extreme inventories

-
22. In discussions with GoC market-makers, several reasons were offered to explain why their desired inventory level changed over time. First, GoC securities are auctioned to dealers on a regular basis, thus causing the inventory levels to fluctuate as dealers take on the securities at auction and then distribute them via secondary market trading. Second, GoC market-makers noted that they often mingle their market-making activities with some minor amount of speculative position-taking. Thus, changes in a dealer's desired inventory positions reflect the changes in the speculative position taken by the dealer and the frequency of primary auctions.
 23. Moreover, GS market-makers do not in general have preferencing arrangements, and therefore must compete strictly on a quote basis for customer market share. This also tends to reduce the bid sharing, since dealers are loath to risk their share of the market-wide order flow for fear of losing their ability to extract information from that flow.
 24. Vitale's results are based on the assumption that desired inventory levels do not vary over time, an assumption that this section argues is unrealistic.

than equity market-makers before engaging in interdealer trading, they nonetheless are relatively active participants in the interdealer sphere, because GS order-flow dynamics are such that these extreme inventory imbalances are frequent occurrences.

Another explanation for the relatively high level of interdealer trading in GS markets is that the dealers' interdealer trading is motivated by information extraction. By trading in the interdealer market, dealers can garner a sense of the market's depth at a given price. This is particularly true for dealers using IDBs, where it has been suggested that dealers engage in price experimentation to extract payoff-relevant information from other dealers.²⁵

A testable implication of this conjecture is that the level of activity observed in the interdealer market bears some relation to the characteristics of the public order flow in the market. Future research could also investigate relationships that exist between interdealer quote activity and the transparency regime in the dealership market (since changes in transparency cause changes in the amount and distribution of information available to market participants). Some related research has already investigated the relationship between interdealer trading activity and transparency (see Scalia and Vacca 1999).

6. Conclusion

This paper has examined the structural differences that exist between multiple-dealer equity and government securities markets. Although multiple-dealer government and equity securities markets have, on the surface, similar market structures, this paper has demonstrated that there are differences between those markets that are likely to significantly affect the way market-makers trade, and as such have an impact on the liquidity that they provide. This paper has also shown that the intrinsic properties of the securities, the trading environment, and customer characteristics in equity and GS markets differ. This fact, combined with the differences in market structure, implies that future research on the functioning of GS markets based on models developed for an equity market framework should be sufficiently adjusted to take the differences between these markets into consideration. Moreover, policy-makers should be wary of applying similar regulations across both equity and GS securities markets, as these rules may in fact be detrimental to the investors that they seek to benefit. Finally, policy-makers should be sensitive to these considerations when applying regulatory prescriptions in one type of market based on lessons or experiences in another.

25. Anecdotal evidence indicates that dealers will sometimes post a firm quote (for a small amount) on an IDB system to see how quickly it is "hit" by other dealers. This gives the dealer posting the quote a sense of the depth underlying such quotes, a sense of the direction and magnitude of the public order flow observed by the other dealers, and/or enables the dealer to extract information that other dealers may possess.

Bibliography

- Allen, H., J. Hawkins, and S. Sato. 2001. "Electronic Trading and Its Implications for Financial Systems." In *BIS Papers No. 7—Electronic Finance: A New Perspective and Challenges*, 30–52. Proceedings of a workshop held by the Bank for International Settlements, July 2001. Available at <<http://www.bis.org/publ/bispap07.htm>>.
- Amihud, Y. and H. Mendelson. 1986. "Asset Pricing and the Bid-Ask Spread." *Journal of Financial Economics* 17: 223–49.
- 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.
- Board, J. and C. Sutcliffe. 1995. "The Effects of Trade Transparency in the London Stock Exchange." Report commissioned by the London Stock Exchange.
- Breedon, F. 1993. "Intraday Price Formation on the London Stock Exchange." London School of Economics Financial Markets Group, Discussion Paper No. 158.
- Cao, H. and R. Lyons. 1998. "Inventory Information." Working Paper, Haas School of Business, U.C. Berkeley, August.
- Committee for Global Financial Stability (CGFS). 1999. "Market Liquidity: Research Findings and Selected Policy Implications." CGFS Report No. 11. Bank for International Settlements, May.
- . 2001. "The Implications of Electronic Trading in Financial Markets." CGFS Report No. 16. Bank for International Settlements, January.
- Copeland, T. and D. Galai. 1983. "Information Effects and the Bid-Ask Spread." *Journal of Finance* 38: 1457–69.
- Dattels, P. 1995. "The Microstructure of Government Securities Markets." IMF Working Paper No. 95/117.
- Fleming, M. and E. Remolona. 1999. "Price Formation and Liquidity in the U.S. Treasury Market: The Response to Public Information." *Journal of Finance* 54(5): 1901–16.
- Flood, M., R. Huisman, K. Keodijk, R. Mahieu, and A. Röell. 1997. "Post-Trade Transparency in Multiple-Dealer Financial Markets." Working paper, Maastricht University.
- . 1999. "Quote Disclosure and Price Discovery in Multiple-Dealer Financial Markets." *Review of Financial Studies* 12(1): 37–59.
- Foster, D. and S. Viswanathan. 1993. "Variations in Trading Volume, Return Volatility, and Trading Costs: Evidence on Recent Price Formation Models." *Journal of Finance* 48: 187–211.
- 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).

- Gemmill, G. 1996. "Transparency and Liquidity: A Study of Block Trades on the London Stock Exchange under Different Publication Rules." *Journal of Finance* 51: 1765–90.
- Glosten, L. and P. Milgrom. 1985. "Bid, Ask, and Transaction Prices in a Specialist Market with Heterogeneously Informed Agents." *Journal of Financial Economics* 14: 71–100.
- Gravelle, T. 1999. "Liquidity of the Government of Canada Securities Market: Stylized Facts and Some Market Microstructure Comparisons to the United States Treasury Market." Bank of Canada Working Paper No. 99–11.
- Grossman, S.J. and M. Miller. 1988. "Liquidity and Market Structure." *Journal of Finance* 43(3): 617–37.
- Hansch, O., N. Naik, and S. Viswanathan. 1998. "Do Inventories Matter in Dealership Markets? Evidence from the London Stock Exchange." *Journal of Finance* 53(5): 1623–56.
- Harris, L. 1993. "Consolidation, Fragmentation, Segmentation and Regulation." *Financial Institutions and Instruments* 2(5): 1–28.
- Hirshleifer, J. 1971. "The Private and Social Value of Information and Reward to Incentive Activity." *American Economic Review* 61: 561–74.
- Ho, T. and H. Stoll. 1983. "The Dynamics of Dealer Markets Under Competition." *Journal of Finance* 38(4): 1053–74.
- Holland, A. 2000. "The Academic Literature—A Brief Review." Annex in DMO Consultation Paper on the Secondary Market for Gilts, January.
- Kyle, A. 1985. "Continuous Auctions and Insider Trading." *Econometrica* 53: 1315–35.
- Leach, J. and A. Madhavan. 1993. "Price Experimentation and Security Market Structure." *Review of Financial Studies* 6: 375–404.
- Lyons, R. 1995. "Tests of Microstructure Hypotheses in the Foreign Exchange Market." *Journal of Financial Economics* 39: 321–51.
- . 1996. "Optimal Transparency in a Dealer Market with an Application to Foreign Exchange." *Journal of Financial Intermediation* 5: 225–54.
- . 2001. *The Microstructure Approach to Exchange Rates*. Cambridge, MA: MIT Press.
- Madhavan, A. 1995. "Consolidation, Fragmentation, and the Disclosure of Trading Information." *Review of Financial Studies* 8: 579–603.
- . 2000. "Market Microstructure: A Survey." *Journal of Financial Markets* 3: 205–58.
- Madhavan, A. and S. Smidt. 1991. "A Bayesian Model of Intraday Specialist Pricing." *Journal of Financial Economics* 30: 99–134.
- . 1993. "An Analysis of Changes in Specialist Inventories and Quotations." *Journal of Finance* 48: 1595–1628.
- Madhavan, A., D. Porter, and D. Weaver. 2002. "Should Securities Markets be More Transparent?" In *Financial Market Structure and Dynamics*. Proceedings of a conference held by the Bank of Canada, November 2001. Forthcoming.

-
- Miyanoya, J., H. Inoue, and H. Higo. 1997. "Market Liquidity of the Japanese Government Securities Market." Bank of Japan working paper.
- Naik, N., A. Neuberger, and S. Viswanathan. 1999. "Trade Disclosure Regulation in Markets with Negotiated Trades." *Review of Financial Studies* 12(4): 873-900.
- O'Hara, M. 1995. *Market Microstructure Theory*. Cambridge: Blackwell Publishers Inc.
- O'Hara, M. and G. Oldfield. 1986. "The Microeconomics of Market Making." *Journal of Financial and Quantitative Analysis* 21: 361-76.
- Pagano, M. and A. Röell. 1996. "Transparency and Liquidity: A Comparison of Auction and Dealer Markets with Informed Trading." *Journal of Finance* 51: 579-611.
- Perraudin, W. and P. Vitale. 1996. "Interdealer Trade and Information Flows in Decentralized Foreign Exchange Market." In *The Microstructure of Foreign Exchange Markets*, edited by J. Frankel et al., 73-106. Chicago: University of Chicago Press.
- Porter, D. and D. Weaver. 1998. "Post-Trade Transparency on Nasdaq's National Market System." *Journal of Financial Economics* 50: 231-52.
- Proudman, J. 1995. "The Microstructure of the U.K. Gilt Market." Bank of England Working Paper No. 38.
- Reiss, P. and I.M. Werner. 1998. "Does Risk Sharing Motivate Interdealer Trading?" *Journal of Finance* 53(5): 1657-1703.
- Saporta, V. 1997. "Which Inter-Dealer Market Prevails? An Analysis of Inter-Dealer Trading in Opaque Markets." Bank of England Working Paper No. 59.
- Scalia, A. and V. Vacca. 1999. "Market Transparency, Liquidity and Price Discovery: A Case Study." Paper prepared for BIS Committee on the Global Financial System Study Group on Market Liquidity.
- Singleton, K. 1996. "Institutional and Regulatory Influences on Price Discovery in Cash and Futures Bond Markets." In *The Industrial Organization and Regulation of the Securities Industry* (NBER conference report), edited by A. Lo, 243-65. Chicago: University of Chicago Press.
- Viswanathan, S. and J.J.D. Wang. 2000. "Inter-Dealer Trading in Financial Markets?" Duke University, Fuqua School of Business Working Paper.
- . 2002. "Market Architecture: Limit-Order Books versus Dealership Markets." *Journal of Financial Markets* 5: 127-67.
- Vitale, P. 1998. "Two Months in the Life of Several Gilt-Edged Market Makers on the London Stock Exchange." *Journal of International Financial Markets, Institutions & Money* 8: 301-26.
- Vogler, K. 1997. "Risk Allocation and Inter-Dealer Trading." *European Economic Review* 41: 417-41.

Appendix A: Literature Review on the Effects of Transparency in Financial Markets

Section A.1 considers theoretical work. Sections A.2 and A.3 discuss empirical and experimental-market studies, respectively.

A.1 Theoretical literature

Pagano and Röell (1996) compare trading costs associated with a dealership market versus a completely transparent order-book market. They show that large traders are better off when trading on the opaque dealership market, while small traders are better off in the transparent order-book market. The drawback of their research is that they assume traders do not care about the trade execution speed (the ability to complete immediate trades). Second, because they don't compare an opaque dealership market with a less-opaque dealership market, but rather compare a dealership market against an order-book market, they are not simply assessing the affects of changing transparency levels, but also assessing the relative usefulness of different market structures.

Madhavan (1995) shows that large institutional traders who do not trade on information related to the fundamental value of the security and who break up their large orders into multiple trades are better off in an environment of less post-trade disclosure, since it reduces their execution costs. That is, large orders can be completed without moving the price in the direction of the trade. Small, uninformed traders (uninformed in terms of the stock's fundamental value), on the other hand, are shown to be better off with greater post-trade information release. Madhavan (2000) states that this research "suggests that one danger of too much transparency is that traders might migrate to other venues" (p. 239). This is particularly relevant for decentralized fixed-income markets in which trading of any debt securities (including Canadian bonds) takes place in many financial centres around the world, since this implies that large liquidity-driven institutional investors could easily migrate their fixed-income trading activity to other, more opaque jurisdictions. That is, market-makers may choose to book their trades in centres with looser transparency regulations—a form of regulatory arbitrage. In fact, an empirical study by Ganley et al. (1998) suggests that Nasdaq market-makers have engaged in regulatory arbitrage by booking large trades they undertook in cross-listed shares on the LSE, where there were no publication requirements, rather than booking the trades on the Nasdaq.

The findings of Naik, Neuberger, and Viswanathan (1999) are of particular interest, since they develop a theoretical model that most resembles dealership markets, which is not the case in the

Madhavan (1995) and Pagano and Röell (1996) studies. The innovation in their approach is that they model the trading that takes place as a two-stage process, in which the dealer first trades with the public investor and then re-trades to manage the inventory position that was taken on in the initial public trade. The key and realistic feature of their model is the fact that the price the dealer offers to the public investor (in the first stage of trading) will depend on the costs of managing its inventory position. They find it essential to model the second stage to understand how prices are set in the first stage. An additional innovation is that Naik, Neuberger, and Viswanathan model both the public investor and the market-maker as being risk-averse and therefore sensitive to price risks. They show that order flow is informative and that dealers will compete for that information by offering preferential quotes. Thus, in contrast with order-book markets, bid-ask spreads do not necessarily increase monotonically with the size of the order, since dealers may rebate the information they gain from the order flow received back to the public investors. This is an important finding for markets that are dominated by large institutions with large orders to transact. The authors also show that an increase in the level of post-trade transparency works against the execution of large trades but nevertheless tends to benefit the small traders.

Lyons (1996) also attempts to model the risk-sharing aspects of dealership markets. He, too, models the realistic feature of two-stage trading in multiple-dealer markets. Lyons presents a model of the FX market (which is very similar to the fixed-income market in terms of trading structure and the characteristics of the market participants). His model allows dealers to off-load their order flow to other dealers in the second stage of trading after having traded with customers in the first stage. Lyons considers the optimal level of transparency in interdealer trading that dealers would choose. (He does not examine the optimal level of transparency for the market as a whole, or consider which level of transparency the investors would choose.) This question is of interest because the FX market is unregulated and yet a degree of transparency has arisen without regulatory prodding.

Lyons shows that dealers, if they could choose the transparency regime *ex ante*, would prefer incomplete transparency over both a completely opaque setting and a completely transparent setting, because incomplete transparency allows for the optimal level of risk-sharing to occur in the market. If there is no transparency, customers do not wish to trade. Under complete transparency, the dealers must bear all the price risk (in the first round) and sharing this risk by trading with market participants in the second round is moot. In essence, by having some positive level of transparency, customer trading activity is sufficient that customers trade with dealers quickly enough (or early enough in the trading session) to allow dealers to lay off some of the price risk on the broader community that includes the customers. If post-trade transparency is complete, prices adjust (nearly) instantaneously and dealers absorb all the price risk

instantaneously. Dealers will therefore prefer some middle ground of transparency. The incomplete degree of post-trade transparency modelled by Lyons is analogous to making the post-trade data from a subset of participants public. This is, in fact, what the publication of IDB post-trade data does (i.e., as in CanPX or GovPX): it provides some degree of incomplete post-trade transparency.

A.2 Empirical literature

Gemmill (1996) examines how changes to the transparency regime on the LSE that occurred in 1989—at that time structured as a dealership market—affected market efficiency and liquidity. Reviewing the data that spanned from 1987 to 1992, Gemmill found no evidence that prices adjusted less quickly (in reaction to news or trading) after a sharp reduction in transparency. His findings agree with those of Breedon (1993), who analyzed a smaller sample of LSE stocks around the same time span. Overall, these results do not support the hypothesis that an increase in post-trade transparency will improve market efficiency. Both studies also found that, on average, relative spreads (measured as the actual price paid compared to the going quote and used to take into account negotiated price improvements for large trades relative to small trades) were narrower under the low-transparency regime. These studies did not reveal any statistically significant effects of decreased post-trade transparency.

Porter and Weaver (1998) lend more support to the idea that excessively stringent disclosure requirements create incentives for dealers to engage in “regulatory arbitrage.” They find empirical evidence indicating that, on average, the trades reported late (or out of sequence) by Nasdaq dealers tend to be large trades and those that are at away prices.¹ The evidence is consistent with the arguments that large trades benefit from delayed reporting.

Madhavan, Porter, and Weaver (2002) also investigate the effects of greater pre-trade transparency using data from the Toronto Stock Exchange (TSE). A regime change in pre-trade transparency occurred on 12 April 1990 when the TSE started to provide real-time public dissemination of the firm quotes for up to four levels away from the inside market (in both directions). Madhavan, Porter, and Weaver find that there is a decrease in liquidity, measured by the realized spread, with the increased transparency regime.

Generally, there is little empirical evidence to support the hypothesis that market efficiency and public investors’ welfare is improved by greater transparency. These studies, however, are carried out as event studies, and therefore changes directly resulting from the changes in transparency

1. An “away price” is defined as a transaction price that differs from quoted or traded prices displayed in the market at the time of the (away) transaction.

regime are difficult to distinguish from other unrelated changes that may have occurred during the sample period in question. Given this fact, it is important to find a “cleaner” environment in which to test the effects of transparency changes. One avenue of research that has recently come along that should help is that of “experimental markets,” in which markets are simulated in a controlled (laboratory) environment.

Another potential drawback of the empirical research is that it has examined transparency regime changes that have occurred on equity markets, which, as stated above, differ in many respects from fixed-income markets. Ignoring for the moment the difference that exists between security and trader characteristics for equity and debt instruments, the empirical research on the LSE or Nasdaq may nonetheless provide some useful insights, since they are structured as dealership markets, just as fixed-income markets are (but differ in that they are centralized markets).

A.3 Experimental literature

Perhaps the research that is most relevant in determining the appropriate level of transparency in a market is that based on experimental economics, in which we can test hypotheses in a controlled environment.

Bloomfield and O'Hara (1999) examine the impact on price discovery and market liquidity of changes in the level of transparency in a multiple-dealer market. They examine three transparency settings: (i) an opaque setting, where customers see neither pre-trade nor post-trade information, (ii) a semi-opaque setting, where customers see the firm quotes posted on a centralized display (a pre-trade transparent setting), and (iii) a transparent setting, which is the same as the semi-opaque one, except that transactions prices and volume are public, as well (i.e., pre-trade and post-trade transparent). They find that price efficiency is not significantly affected when going from opaque to semi-opaque, but find that prices move rapidly to their new fundamental level (when new information about the fundamentals is released) in the transparent setting. They also find that changes in transparency cause significant change in the traders' profits. As with most research in this area, by separating the traders into three groups of informed traders, uninformed or liquidity-motivated traders, and dealers, they show that informed traders are made worse off by going to the most transparent setting. This is not surprising, since it is predicted in the theoretical work: traders can no longer profit from their insider information when trading is transparent. There was no significant change in the traders' returns when going from the opaque to the semi-opaque setting. The liquidity traders, however, who need immediacy, are made significantly worse off in the transparent setting. This is surprising, since it contradicts the findings of Pagano and Röell (1996), who predict that an increase in transparency would shift the distribution of returns from the

informed traders towards the uninformed liquidity traders. They also find that dealers compete less for order flow and that spreads are wider in the transparent setting.

Flood et al. (1999) conduct an experimental economic study modelled as a multiple-dealership market. Their work differs in that traders face a time pressure that places a greater emphasis on the trader's learning ability. They also examine only changes in pre-trade transparency, leaving the post-trade environment opaque. They find that as the level of pre-trade transparency increases, trading activity increases, although the returns of the traders are not significantly affected in a statistical way. Surprisingly, however, they find that greater pre-trade transparency slows down this price-discovery process (i.e., price efficiency is significantly decreased). This is in marked contrast with Bloomfield and O'Hara (1999) and with much of the theoretical literature on transparency. Flood et al. hypothesize that this difference occurs because dealers, no longer needing to seek out pre-trade information by calling for and making quotes, become less aggressive price setters in the more transparent setting, and therefore the price converges more slowly to the one predicated by fundamentals. In essence, the greater transparency reduces the incentives of the dealers to acquire new information (which in turn is revealed/embodied in the price).

Although both Flood et al. (1999) and Bloomfield and O'Hara (1999) replicate many of the features of a dealership market, one feature they do not replicate is the dealers' ability to trade amongst themselves via IDBs. The dealers' exclusive access to this platform and the fact that IDBs tend to provide a partial signal of the market-wide order flow could play a vital role in increasing the price efficiency of the market. The author of the present paper hypothesizes that any additional transparency imposed on a multiple-dealer market that has IDB systems would likely yield little price-efficiency improvements.

Bloomfield and O'Hara (2000) shed some light on regulatory arbitrage questions. In their paper, dealers can either disclose their trades (thus being post-trade transparent) or not disclose them. The authors find that non-transparent dealers are more aggressive in setting quotes and that this allows them to capture more order flow. This greater share of order flow provides dealers with an informational advantage (over the disclosing dealers) and allows them to quote narrower spreads. Although not directly examined in the study, this in principle benefits the non-dealer traders. In the end, the low-transparency dealers' profits are significantly higher than the transparent dealers' profits. Bloomfield and O'Hara suggest that markets will naturally gravitate towards less-transparent regulatory settings.

In each of the experimental studies discussed above, informed traders are assumed to be active. In fixed-income markets, it is hard to imagine that an investor would have superior (insider)

information about the value of the security; it is likely that this type of trader does not exist or comprises a minority of investors. Thus, the gains and losses of the uninformed liquidity trader owing to changes in transparency would be the prevailing consideration.

A.4 Summary

It seems that, in terms of transparency, no optimal market structure exists that would benefit all types of market participants or accommodate, at the same time, increased market efficiency and liquidity. The current state of research calls into question the view that increased transparency is unambiguously beneficial for markets, and suggests that market structure and investor attributes should be important considerations in assessing the appropriate transparency regime.

Appendix B: The Effects of Greater Post-Trade Transparency in Dealership Markets

An economic trade-off arises with greater price transparency. Although greater transparency can accelerate the embodiment of information into prices, it can also impede dealer risk management. That is, although greater transparency can make markets more informationally efficient, greater transparency can also reduce the ability of dealers to perform inventory risk-sharing or risk management, which makes the trading process less efficient.² This appendix examines this trade-off in greater detail.³

The tension between the benefits and costs of greater transparency would not exist if perfect inventory risk-sharing could actually be accomplished. In a market where there exist K market-makers, perfect inventory risk management is accomplished for a customer order of size X when there is *perfect diversification*, where the order is equally divided among the K market-makers (each holding X/K of the order), and where risk-sharing is *instantaneous*. Inventory risk-sharing, however, rarely achieves perfect diversification, nor does it take place instantaneously.⁴ If *instantaneous* risk-sharing could be accomplished, greater transparency would likely be beneficial, irrespective of the timing of the information release (or the size of the customer order).

Given that dealers cannot predict their inventory levels before customer orders arrive, dealers face unavoidable inventory disturbances or imbalances, because they maintain a continuous presence in the market. Thus, dealers seek to manage the risks that arise from these inventory imbalances by conducting interdealer trading. Greater post-trade transparency—defined here as the imposition of immediate publication of the price and size of completed transactions between an investor (customer) and a dealer—accelerates the revelation of information in prices, increasing the speed at which prices move away from the dealer's executed or quoted price, and/or their probability of doing so, leaving dealers less time to manage their inventory imbalances. The nearer the disclosure of post-trade information is made to the time at which the transaction actually takes place, the more the dealer's inventory risk-management costs increase. This in turn

-
2. Greater transparency *might not* accelerate information revelation in the price if one takes into account the fact that greater transparency reduces incentives to acquire new information.
 3. The issues related to asymmetric information are ignored in the discussion. They would likely complicate some of the inferences made in this appendix and they are at the forefront of market microstructure research.
 4. The efficiency gain achieved from inventory risk-sharing (i.e., the diversification gain) depends on the time it takes for interdealer trading to achieve such inventory risk-sharing. The longer it takes for interdealer trading to achieve inventory risk-sharing, or the more imperfect the inventory risk-sharing, the smaller the efficiency gains.

might also cause the dealer to conduct less inventory risk management. Thus, greater post-trade transparency decreases incentives for interdealer trading and reduces socially efficient risk-sharing. Moreover, these greater risk-management costs—which lead to lower market liquidity and impede the price discovery—are passed on to customers, who are worse off.

How does greater transparency make it more difficult and more costly for dealers to carry out inventory risk management? Recall that a dealer trading in the interdealer market must pay other dealers a bid-ask spread, just as a customer would, to acquire immediate trade-execution services. As greater transparency increases the inventory risks faced by (all) dealers, risk-averse dealers will naturally quote wider bid-ask spreads in the interdealer market (in tandem with the greater risks they face). These wider spreads increase the cost of inventory risk management, leading the initiating dealer to undertake less of it. (In a competitive environment, these higher risk-management costs are passed on to the customer.)

Hirshleifer's (1971) work shows that another way to think about inventory risk management is to consider dealers as being subject to endowment risks. He shows that, in general, additional ex ante information can impede risk-sharing. Assuming that risk-averse agents are endowed with n state-contingent goods (or n goods that represent the n states), and given a set of state probabilities, Hirshleifer shows that if information about the true state of the world is revealed before trading can occur or be completed, optimal endowment risk (re)allocation is not achieved. Using this framework, one can understand how greater post-trade transparency increases the endowment (or inventory) risk faced by dealers in conducting interdealer trading. That is, as dealers are continually hit by inventory (or endowment) disturbances, they will seek to reallocate their inventory positions. Because there is a time dimension to this process, when price revelation is too quick, inventory risk management is inhibited. Stated differently, immediate post-trade transparency decreases the efficient reallocation of these risks.

Some delay in the disclosure of post-trade information will often be sufficient to mitigate its detrimental impact on inventory risk management. Delaying the release to the public of completed dealer-to-customer trade information would allow for the efficient reallocation of inventory imbalances, and would likely not negatively impact market liquidity.

If customer orders are large on average, inventory disturbances will be larger on average, which will accentuate the dealer's risk-management needs. Moreover, there is a greater time dimension to larger orders. The larger the inventory disturbance or customer order, the longer it takes to conduct inventory risk-sharing. These factors (size and time) accentuate a dealer's sensitivity to accelerated information revelation. At some point, as the average size of the customer orders increase, the benefits of real-time post-trade transparency are dominated by the costs arising from

the decrease in socially beneficial inventory risk-sharing. When the (average) customer order is small, the benefits of greater transparency likely outweigh the costs (from having less inventory risk-sharing). But for larger orders, the gains provided by inventory risk-sharing are more important. Given the time dimension to inventory risk management, larger orders require a greater delay in post-trade information disclosure (i.e., less trade transparency), to mitigate the detrimental effects of price revelation on a dealer's inventory risk-management ability.

Thus, in certain markets, where the average trade size is relatively small, as in the case of equity dealership markets (e.g., Nasdaq), the imposition of real-time post-trade transparency might improve customer welfare on average. But in markets where the average trade size is relatively large and where customer orders are relatively infrequent, as in the case of fixed-income markets, real-time post-trade transparency is likely to reduce the amount of liquidity available for investors (i.e., it increases the costs of immediacy for customers), and is likely to reduce market efficiency and investor welfare.⁵

To summarize, greater or enhanced transparency, such as real-time post-trade information disclosure, can readily reduce inventory-risk diversification in certain dealership markets (like fixed-income markets) and so reduce market liquidity and ultimately investor welfare.

5. Most of the academic literature that examines the effects of transparency in dealership markets is quite recent, and the results from that literature, while insightful, should be considered to be tentative (see Appendix A).

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

- | | | |
|--------|---|------------------------------------|
| 2002-8 | Restructuring in the Canadian Economy: A Survey of Firms | C. Kwan |
| 2002-7 | Contribution of ICT Use to Output and Labour-Productivity Growth in Canada | H. Khan and M. Santos |
| 2002-6 | Currency Fluctuations, Liability Dollarization, and the Choice of Exchange Rate Regimes in Emerging Markets | P.N. Osakwe |
| 2002-5 | The Effects of Bank Consolidation on Risk Capital Allocation and Market Liquidity | C. D'Souza and A. Lai |
| 2002-4 | Does Micro Evidence Support the Wage Phillips Curve in Canada? | J. Farès |
| 2002-3 | An Introduction to Wavelets for Economists | C. Schleicher |
| 2002-2 | Asset Allocation Using Extreme Value Theory | Y. Bensalah |
| 2002-1 | Taylor Rules in the Quarterly Projection Model | J. Armour, B. Fung, and D. Maclean |

2001

- | | | |
|---------|---|--------------------------|
| 2001-27 | The Monetary Transmission Mechanism at the Sectoral Level | J. Farès and G. Srouf |
| 2001-26 | An Estimated Canadian DSGE Model with Nominal and Real Rigidities | A. Dib |
| 2001-25 | New Phillips Curve with Alternative Marginal Cost Measures for Canada, the United States, and the Euro Area | E. Gagnon and H. Khan |
| 2001-24 | Price-Level versus Inflation Targeting in a Small Open Economy | G. Srouf |
| 2001-23 | Modelling Mortgage Rate Changes with a Smooth Transition Error-Correction Model | Y. Liu |
| 2001-22 | On Inflation and the Persistence of Shocks to Output | M. Kichian and R. Luger |
| 2001-21 | A Consistent Bootstrap Test for Conditional Density Functions with Time-Dependent Data | F. Li and G. Tkacz |
| 2001-20 | The Resolution of International Financial Crises: Private Finance and Public Funds | A. Haldane and M. Kruger |

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>