



American Finance Association

Does Risk Sharing Motivate Interdealer Trading?

Author(s): Peter C. Reiss and Ingrid M. Werner

Source: *The Journal of Finance*, Vol. 53, No. 5 (Oct., 1998), pp. 1657-1703

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/117420>

Accessed: 25/11/2009 09:59

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=black>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing and American Finance Association are collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Finance*.

<http://www.jstor.org>

Does Risk Sharing Motivate Interdealer Trading?

PETER C. REISS and INGRID M. WERNER*

ABSTRACT

We use unique data from the London Stock Exchange to test whether interdealer trade facilitates inventory risk sharing among dealers. We develop a methodology that focuses on periods of “extreme” inventories—inventory cycles. We further distinguish between inventory cycles that are unanticipated and those that are anticipated because of “worked” orders. The pattern of interdealer trade during inventory cycles matches theoretical predictions for the direction of trade and the inventories of trade counterparts. We also show that London dealers receive higher trading revenues for taking larger positions.

IN DEALER MARKETS, TWO OR MORE broker-dealer firms compete for customer transactions by quoting bid and ask prices, and sizes.¹ Prominent examples of dealer markets include Nasdaq, London’s Stock Exchange Automated Quotations System (SEAQ), and spot foreign exchange markets. Several empirical studies suggest that interdealer trade accounts for a substantial fraction of all trade in these markets. For example, Reiss and Werner (1995) estimate that interdealer trading on the London Stock Exchange accounts for 25

*Both authors are from Graduate School of Business, Stanford University, and the National Bureau of Economic Research. We thank Lee-Bath Nelson, Mark Garmaise, and Mary Pressley Vyas for expert research assistance. We obtained the data for this project from the Quality of Markets Group at the London Stock Exchange. We especially thank Stephen Wells and Graham Hart for helping us interpret the data. We also are grateful to George Gray of Union Bank of Switzerland, Richard O’Hare of Salomon Brothers, Xavier Rolet of Goldman Sachs & Co., David Bamba, Colin Taylor, and Jan G. Werner of Merrill Lynch, Bradley Bilgore and Ashley Short of Morgan Stanley & Co., Leslie O’Malley and Rod Williams of Tullett & Tokyo, and Peter Barnett of Tradepoint for answering our many questions about the London market. Feedback on our presentations at Cornell University, the Hoover Institution, the Institute for Fiscal Studies, Indiana University, the Massachusetts Institute of Technology, the New York Stock Exchange, Strathclyde University, the University of Chicago, the University of Southern California, Vanderbilt University, Washington University, The University of Memphis Competition for Order-flow Conference (1995), The Ohio State Market Dealer Markets Conference (1996), and the London School of Economics Market Microstructure Conference (1997) helped us tremendously. We also are grateful to Narayan Naik, René Stulz, Sunil Wahal, and an anonymous referee for comments that greatly improved the paper. This paper was completed while Werner was a Visiting Research Economist at the New York Stock Exchange. A Hoover Institution National Fellowship also supported Werner’s research. Both authors received support from Stanford’s Financial Research Initiative. The usual disclaimer applies.

¹ In London, market makers, or in our terminology “dealers,” are called broker-dealers because they can represent agency orders as brokers and serve as market makers as well. By size we mean the maximum amount that dealers agree to trade at their bid and ask prices.

to 35 percent of total share turnover. Gould and Kleidon (1994) estimate from a small sample of Nasdaq trades that interdealer volume represents roughly 15 percent of turnover on Nasdaq. Lyons (1996) estimates that interdealer trades account for 85 percent of turnover in spot foreign exchange trading.

Despite the prevalence of interdealer trade, relatively few academics have studied it. Existing theoretical models have explored the role that interdealer trades might play in facilitating risk sharing. As O'Hara (1995) notes in commenting on Ho and Stoll's (1983) model, theoretical research into interdealer trade has been slowed by the difficulty of working with models in which dealers are asymmetrically informed. Empiricists also have been stymied in their effort to study interdealer trade. The primary obstacle to empirical research is the lack of publicly available intraday transactions data that identify interdealer trades.

This paper uses unique data from the London Stock Exchange (LSE) to study the extent, causes, and consequences of interdealer trade. Our main hypothesis is that dealers primarily use interdealer trades to reduce their inventory exposure. Though this hypothesis may seem somewhat obvious, it often is lost in debates about the efficiency of dealer markets. We begin by documenting the extent of interdealer trading and the degree to which it is associated with inventory imbalances resulting from customer and broker trades. Guided by the inventory risk-sharing models of Ho and Stoll (1983) and others, we examine both whether interdealer trades reduce the inventories of trading counterparties and whether interdealer trades increase during inventory imbalances. To support our hypothesis that it is primarily inventory risk that motivates dealers to trade with one another, we also test whether dealers earn greater trading revenues in return for assuming greater inventory risk. Finally, we examine whether interdealer trades themselves affect the trading revenues dealers receive when assuming inventory risk.

There are several features of the London equity markets and our London intraday transaction data that facilitate our statistical tests. First, we can identify interdealer trades and trade counterparties. Second, because Exchange regulations require dealers to accept large trades at their quoted prices, London dealers may quickly acquire significant positions involuntarily. Third, dealers often take block trades voluntarily, acquiring million-pound (or more) long or short positions in a few trades. These large trades create considerable inventory risk, some of which is not cheaply hedged. It might seem as though dealers could move their quotes to attract customer order flow and thereby reduce their inventory risk, but other researchers have shown that preferencing arrangements and adverse selection risks can make this a costly strategy.

To our knowledge, our work is the first comprehensive empirical study of interdealer trade in equity markets. It is related to studies of specialists' and derivatives traders' inventory policies (e.g., Madhavan and Smidt (1993), Hasbrouck and Sofianos (1993), and Manaster and Mann (1996)). The closest paper to ours is one by Hansch, Naik, and Viswanathan (1998) that

generally discusses the inventory management policies of London dealers. Our analysis differs from previous empirical studies of dealers' inventory policies in several ways. First, we distinguish between inventory positions resulting from customer, broker, and other types of trades. Second, we distinguish between inventory imbalances resulting from anticipated and unanticipated customer and broker trades. According to the Exchange, roughly 30 percent of large customer trades are anticipated orders that dealers "work" during a day.² We find that failure to account for these anticipated trades seriously compromises our tests. Third, we examine whether risk-sharing motivates *both* parties participating in the interdealer trade. Fourth, unlike most previous inventory management studies, we use intraday trade data to construct dealers' intraday inventories. Using these inventory data, we then identify "inventory cycles" in which dealers' inventories depart from a normal level, reach an extreme, and then return to a normal level. We focus on intraday inventories because most inventory cycles in London are short-lived. Finally, our analysis exploits variation in the execution costs of interdealer trades to explain differences in dealers' risk-sharing strategies (as represented by the mix of interdealer trades) and differences in dealers' trading revenues.

Our main findings are as follows. We report results for two groups of stocks: twenty-five Financial Times-Stock Exchange 100 index stocks (FTSE stocks), and twenty less liquid stocks (non-FTSE stocks). Interdealer trading accounts for 24 percent (by value) of all equity trading. It is evenly split between anonymous electronic trades and direct phone-negotiated trades. When a dealer has a neutral inventory position, interdealer trades account for less than 10 percent of the dealer's trades. When a dealer's inventory reaches what we define as an "extreme" level in an inventory cycle, either because of one large trade or several medium-sized customer trades, the frequency of interdealer trading increases to between 55 percent (non-FTSE stocks) and 65 percent (FTSE stocks) of all trades.³ During inventory cycles associated with customer trades, a high proportion of the interdealer trades made by a dealer, roughly 80 percent, reduce that dealer's current (or anticipated) inventory imbalance. Furthermore, more than 60 percent of all interdealer trades occur between dealers who are *simultaneously* reducing inventory imbalances. Additionally, we find that interdealer counterparties frequently have the most extreme (and opposite) inventory positions.

² See the LSE consultation document, *New Electronic Trading Service*. Other studies that explore this distinction include Board and Sutcliffe (1995) and Oliver, Wyman & Co. (1997).

³ We attribute 86 to 95 percent of the inventory cycles in our data to customer or broker trades. When we analyze only those cycles caused by broker trades, interdealer trading increases to 80 percent (FTSE stocks) to 90 percent (non-FTSE stocks) of the trading volume that reverses the inventory imbalance. For the 5 to 14 percent of all inventory cycles in which we cannot unambiguously attribute the cycle to customer or broker trades, we find that interdealer trading accounts for 65 percent (non-FTSE stocks) to 70 percent (FTSE stocks) of the trading volume that reverses the inventory imbalance.

London dealers also actively manage their extreme inventory positions (although not as actively as the derivatives traders studied by Manaster and Mann (1996)). More than half of all FTSE stock inventory imbalances caused by customer or broker trades are reversed in a day. By contrast, the cycles that we cannot classify as related to customer or broker trades last on average seven days. Finally, we find that dealers who take on greater inventory risks are rewarded for this risk taking with greater trading revenues. Taken together, these results suggest that London dealers are risk averse and that they primarily use interdealer trades to share inventory risks.

Section I describes the London market, different types of interdealer trades, and the extent of interdealer trading. Section II outlines five testable hypotheses motivated by the theoretical literature on interdealer trading. Section III describes how we implement our tests, including how we define dealer inventory cycles. Section IV reports our tests for a sample of twenty-five FTSE stocks, and twenty less liquid non-FTSE stocks traded during 1991.

I. London Interdealer Trading

Since 1986, the London Stock Exchange has operated as a dealer market. Although it is similar to Nasdaq, there are important differences (see, e.g., Neuberger and Schwartz (1989) and Tonks and Webb (1989)). Perhaps the most important difference is that the Exchange requires dealers to accept large trades. The mandated quote size, referred to as the normal market size (NMS), equals 2.5 percent of a stock's prior quarter average daily trading volume. During 1991, the period of our study, most FTSE stocks had NMSs that required them to accept trades of £250,000 to £500,000 if they were posting at the best bid or offer. The large mandated sizes in London reflect the desire of many institutional traders to execute large orders or "blocks". According to the Exchange, during 1991 trades exceeding £1,000,000 accounted for more than one-third of the total turnover in large stocks (*Stock Exchange Quarterly* (1991)). These statistics suggest that dealers voluntarily and involuntarily acquire substantial inventory in short periods of time.

A common misperception is that London dealers can manage their inventories effectively simply by moving to, or improving on, the best bid or offer. In such a world, interdealer trading would be unnecessary. Research by Hansch et al. (1998) has shown, however, that London dealers attract little additional customer order flow by moving their public quotes. The reason is that more than 70 percent of customer order flow is preferenced (Hansch, Naik, and Viswanathan (1997)). This finding suggests one reason why dealers might use interdealer trades to manage their inventories. Indeed, in London there are four different ways dealers and brokers can trade with one another. One popular way is to use one of four interdealer broker (IDB) services. The IDBs provide anonymous electronic matching services similar to Reuters' Instinet. A second way is by direct phone negotiation, a practice we will refer to as IMM (intermarketmaker) trading.

Table I
1991 London Trade Volume by Trade Type

We construct this table using individual trade data supplied by the Exchange's settlement system (see Appendix A). IDB refers to trades consummated on one of the four interdealer broker systems. IMM refers to direct trades between dealers registered to quote prices for a security. Broker refers to proprietary trades by Exchange brokers who are not dealers. Broker trades include crosses and trades between a broker and a dealer. Customer trades are trades where at least one counterparty is someone other than a London dealer or broker.

Trader Type	U.K. and Irish Equities		FTSE Equities	
	Volume (million £)	Volume (%)	Volume (million £)	Volume (%)
Interdealer				
IDB	24,948	9.2	19,453	10.7
IMM	24,936	9.2	20,154	11.1
Broker	13,647	5.0	10,339	5.7
Total	63,531	23.5	49,946	27.5
Customer	206,655	76.5	131,664	72.5
All trades	270,186	100.0	181,610	100.0

The last two types of interdealer trades involve brokers who are not registered dealers in the stock: here a registered dealer trades with a broker or two brokers execute a cross.

Table I summarizes the extent of London interdealer trading in U.K. and Irish equities during 1991, which is our sample period. Appendix A describes our data in more detail, its coverage and the procedures used to construct Table I. During 1991, London trading volume totaled roughly two-thirds of the trading volume on Nasdaq, or about £270 billion. By value, interdealer trades accounted for 23.5 percent of all trades.⁴ Table I reveals that IDB and IMM trades were equally important, with each representing 9.2 percent of total turnover value. Trades between two brokers or a dealer and a broker not registered in the stock accounted for the remaining 5.0 percent of trades. The second set of columns reports trading volume for 100 of the largest U.K. and Irish equities (the FTSE). FTSE stocks accounted for just over two-thirds of the total turnover in U.K. and Irish equities. Interdealer trades are relatively more important for the liquid FTSEs, accounting for 27.5 percent of turnover.⁵

⁴ These numbers are comparable to those reported in Reiss and Werner (1995), except the fraction of IDB trades is almost halved. The decrease occurs because in that paper we adopt the Exchange's convention of double counting IDB trades: the Exchange records an IDB trade as both a buy from an IDB and a sell to an IDB. Here we count IDB trades as one trade between two market makers.

⁵ The London Stock Exchange separately tabulates interdealer trade. It estimates that interdealer trading accounts for 35 to 50 percent of all trade, well above what we report. We attribute much of the difference to the Exchange's double counting of interdealer trades. Small differences come from differences in data. Table AI in Appendix A compares our numbers to

In Reiss and Werner (1995,1997a) we describe the institutions supporting interdealer trades. We also analyze the relative costs of IDB, IMM, trades between brokers, and between brokers and dealers. In these papers, we show that a dealer who calls a competitor to negotiate a principal trade rarely obtains price improvement relative to the posted quotes. By contrast, if a dealer hits a competitor's limit order on an IDB screen, he improves on the public quotes, on average, by one-third of the public market best bid-ask spread (called the "touch" in London). Finally, the typical broker-to-broker or broker-to-dealer trade has the counterparties evenly splitting the touch. In Reiss and Werner (1995), we also show that although small to medium-sized customer trades in London typically do not obtain any price improvement, many large customer trades do take place inside the touch.

Spread revenue, or more generally trading revenue, is but one component of dealer revenues. If a dealer is acting as a broker representing a customer and takes the other side on his own book, he also may charge commission. According to a recent study by Hansch et al. (1997) using 1994 data, roughly 60 percent of customer trades in London are "internalized" in this way. Unfortunately, we do not have access to the trade commission data necessary to compare the revenues dealers receive from different types of trades. The Exchange's annual transactions surveys, however, reveal that brokerage commissions (in basis points) decline with the size of a customer trade (*Stock Exchange Quarterly* (1991)). Further, 25 percent of customer trades by value are free of commissions. Thus, a dealer charging average commission rates and with a normal order flow would earn net commissions ranging from more than 100 basis points for customer trades below £20,000 to 20 basis points for customer trades exceeding £1,000,000. By contrast, dealers do not receive commissions when they execute interdealer trades.

Taken together, these calculations suggest that London dealers should prefer to use small to medium-sized customer trades to unwind inventory imbalances. Small to medium customer trades pay the full spread and generate more commission than either large customer trades or interdealer trades. The problem for dealers of course is that there might be insufficient customer business that would allow them to unwind an imbalance in a timely manner. Should a dealer have to resort to interdealer trades, we would expect him to first try the anonymous IDB screens. Only as a last resort would he pick up the phone and call a competing dealer to negotiate a trade. At this juncture, we note that the sheer volume of interdealer trading in Table I suggests that customer liquidity is typically insufficient to provide dealers with opportunities to manage large inventory positions. Furthermore, dealers find it necessary to use the apparently costly phone-negotiated interdealer trades.

those published by the Quality of Markets Group. Our estimated fraction of interdealer trading is more in line with estimates of interdealing trading from other equity markets such as Nasdaq. Using separate London data for January 1995, we find interdealer trading accounted for 27.4 percent of volume in all U.K. and Irish equities. Of this 27.4 percent, IDB trades accounted for 9.5 percent, IMM 12.5 percent, and broker 5.4 percent.

II. Hypotheses

Despite the importance of dealer markets for equity trading, relatively few theoretical models explicitly model interdealer trade. The first paper to model interdealer trading appears to be that of Ho and Stoll (1983). They analyze a dealer market where risk-averse dealers use bid and ask quotes to compete for liquidity-motivated customer orders. Customer trades, and thus dealer inventories, are public information.⁶ Ho and Stoll's main focus is to relate market spreads to the distribution of inventories among dealers, but they also discuss the incentives dealers have to engage in interdealer trade.

More recent models of interdealer trade include Vayanos's (1998) dynamic linear rational expectations model, Vogler's (1993) static model, and Naik, Neuberger, and Viswanathan's (1996) model.⁷ These papers study settings where risk-averse dealers compete for a large trade that they then can lay off by trading strategically with other speculators. All three models share the feature that interdealer trading allows the dealer who took the trade to share inventory risk with other dealers. In Naik et al. (1996), the market maker taking the trade has the additional benefit of being able to exploit any information contained in the order when trading with other (less informed) dealers.⁸ In practice, dealers may also obtain information from sources other than trades. Regardless of the source of information, dealers may use interdealer trades when responding to the information. We have little to say about how frequently dealers use interdealer trades in this way, mainly because we do not know when dealers have private information. In the rest of this section, we outline five hypotheses about interdealer trading based on the idea that dealers primarily use interdealer trades to reduce inventory risks.

Each of the above-described theoretical models suggests that inventory risk sharing is a principal benefit of interdealer trade. These benefits do not come without cost, the main cost being that the initiating dealer must give up all or part of the spread when dealing with a competitor. As O'Hara (1995) notes, although models of interdealer trade contain special assumptions and cannot yet be solved in dynamic settings, the basic intuition is that the benefits from risk sharing are more likely to exceed the costs when dealers are extremely far from their desired inventory position. Just how large the benefits are and how much trade occurs depends on the equilibrium bid and ask prices set by other dealers. In Ho and Stoll's (1983) duopoly model, for example, no interdealer trade occurs because the dealers set prices

⁶ Biais (1993) studies the more realistic case where inventories are private information. He does not, however, allow interdealer trade. See also Werner (1997) for a sealed-bid double auction model of interdealer trading.

⁷ These models are extensions of Kyle (1989).

⁸ Naik et al. (1996) are inspired by Exchange rules that delay the public disclosure of large customer trades. The definition of what constitutes a large trade in the context of delayed disclosure has varied over time. During 1991, a large trade was any customer trade in excess of three NMS and it was automatically delayed for 90 minutes. Our data report actual trade times, not when trades are made public.

that make it profitable to wait for customer orders. Ho and Stoll claim that as the number of dealers increases, interdealer trades occur and, further, the volume of interdealer trade increases with an increased dispersion in inventories. Because their model does not provide specific predictions about the extent of interdealer trade, we take as our null hypothesis the opposite claim that the frequency of interdealer trading is unrelated to dealers' individual inventory positions or the joint distribution of inventory positions. Our alternative hypothesis is the one suggested by Ho and Stoll:

HYPOTHESIS 1: Interdealer trades should increase when a dealer's inventory position is "extreme."

If dealers are particularly sensitive to inventory risk, one might expect them to immediately use interdealer trades to reduce that risk. This possibility does not arise in Ho and Stoll's model because they assume that dealers have fixed trade sizes and fixed trade opportunities. Thus, in their model, dealers will trade in the direction of the mean market inventory, but typically do not eliminate all inventory risk. Though the size of interdealer trades in London is not preset, there are several reasons why dealers might not use interdealer trades to unwind positions instantaneously. First, dealers find continuous inventory management prohibitively expensive, especially because a typical dealer manages 20 or more stocks simultaneously. Second, both IDB and IMM trades by convention tend to occur in integer multiples of NMS units, with one NMS being the modal size (Reiss and Werner (1997a)). Third, when dealers try to trade in large sizes, they may adversely impact the price. Thus, although Ho and Stoll's (1983) assumptions do not literally fit the London interdealer trading environment, we test the null hypothesis that the direction of interdealer trades is independent of dealer inventories against the alternative:

HYPOTHESIS 2: When a dealer's position is "extreme," interdealer trades should be in the direction of a neutral or "desired" inventory position.

Ho and Stoll's model makes even stronger predictions about the pattern of interdealer trade and how it relates to the distribution of inventories among dealers.⁹ The most competitive quotes would be posted by the dealers at extremes of the inventory distribution. If dealers traded sequentially, we would expect the first trade to be between the dealer with the lowest inventory (the buyer) and the dealer with the highest inventory (the seller). The next trade would be between a potentially new pair of dealers with extreme inventories, and so on. We take as a null hypothesis that there is no relationship between where dealers are in the inventory distribution versus the alternative hypothesis:

HYPOTHESIS 3: When interdealer trades do occur, they should take place between the two dealers with the most "extreme" inventory positions.

⁹ By inventories we mean the extent to which a dealer's actual inventory deviates from that dealer's desired inventory position.

We examine Hypotheses 1 through 3 using information on dealers' inventories surrounding large customer and broker trades, or sequences of customer and broker trades. We focus on positions caused by nonproprietary trades because they are less likely to be related to instances when a dealer is actively taking a position in the stock (speculation). For nonspeculative positions, we find it reasonable to assume that the desired long-run inventory level is zero. Note that this is not to say that our hypotheses do not also apply to speculative positions. However it is impossible, with the available data to estimate the desired inventory level for a purely speculative position.

Though the Exchange requires dealers to accept orders of one NMS or less at their quoted prices, dealers also frequently accept larger orders. Reiss and Werner (1995) show that block traders often obtain price improvement. Exchange survey data also show that these large trades pay low or no commissions.¹⁰ Thus, unless dealers profit from block orders through subsequent trades, it appears that dealers receive little if any premium for assuming added inventory risk. Similarly, we would expect risk-averse dealers to require compensation for taking speculative positions. We take as our null hypothesis that dealers are risk neutral and that they are not in a position to price discriminate based on order size. Our alternative hypothesis is:

HYPOTHESIS 4: Risk-averse dealers should receive a risk premium for allowing their inventories to deviate substantially from neutral or "desired" inventory positions.

Our previous empirical findings (Reiss and Werner (1995,1997a)) on the spread costs of different types of trades suggest that interdealer trades should be less profitable than customer trades. Of course, the instantaneous spreads on customer and interdealer trades do not capture differences in the timing of these trades relative to quote changes. Thus, an analysis of the profitability of interdealer trading must consider whether interdealer trades are timed in a way that affects their profitability. Since interdealer trades increase around the trading of large customer trades, we must separate the effect of the block on trading revenues from that of the interdealer trades. Our null hypothesis is that London dealers earn the same amount of trading revenues on interdealer trades as they do on customer trades. Our alternative hypothesis is:

HYPOTHESIS 5: When unwinding extreme inventory positions caused by customer trades, dealers earn more trading revenues from customer trades than they do from interdealer trades.

To test these five hypotheses, we need to know or estimate individual dealers' inventories. This is the topic to which we now turn.

¹⁰ We note that this pattern contradicts some market microstructure models that suggest large trades should receive worse execution (see, e.g., Easley and O'Hara (1987)).

III. Dealer Inventories

A unique feature of our transaction data is that they contain settlement information that enables us to estimate dealer inventories. Specifically, we know when a dealer is taking a trade as a principal (for his own account) and whether he is buying or selling shares. An estimated inventory series for a dealer can be obtained by cumulating the trades he or she takes as a principal over the course of our sample year, 1991. In principle, this can be done for all dealers for every one of the 1,923 stocks in our sample. In the following subsections we discuss potential pitfalls associated with inventory calculations, give examples of the diversity of dealers' inventory management strategies, introduce a method for categorizing the inventory series into "extreme" positions and "desired" ranges of inventories, and explain why distinguishing between anticipated and unanticipated customer trades is important.

A. Measurement Issues

There are several reasons why inventories obtained by cumulating trades may have little to do with a dealer's exposure to inventory risk. Consider, for example, the case where we cumulate the trades of dealers in a London equity that also trades overseas. In this case, a dealer's London position may have little to do with his worldwide position. Consider also the case of a dealer who trades an equity that has warrants, options, or futures. Here again the dealer's position in the underlying equity need bear no relation to the capital at risk. To avoid these problems, we exclude from our analyses any stock that has significant trading activity abroad. We also eliminate any stock that had separate rights issues or options trading during 1991.¹¹ These criteria, and the criterion that we have complete transaction data, reduce our candidate sample equities to 25 FTSE stocks and several hundred non-FTSE stocks. To limit our calculations, we work with only 20 non-FTSE securities. We choose these securities so that they represent a range of market capitalizations and dealers. Specifically, we rank the candidate non-FTSE stocks first by the maximum number of competing dealers they had in 1991, and then by their March 1991 market capitalization. For a given number of dealers (which ranges from 4 to 13), we then choose to include the two stocks with the largest market capitalizations. Table BI in Appendix B lists these stocks and provides summary statistics.

To construct inventories, we cumulate all trades by the dealers on their own account by stock for all of 1991. We are careful to exclude the between 1 and 2 percent of trades that dealers canceled because of reporting errors or mistakes. Because we do not know the inventories that the dealers had at

¹¹ Appendix B provides more details on how we select stocks. Ideally, we should also evaluate the opportunities dealers have to hedge their inventory positions with substitute financial instruments.

the start of 1991, we difference each dealer's inventory for each stock from the annual average of that dealer's daily closing inventory.¹² We use this adjusted inventory series in our calculations.

Figure 1 displays four inventory series for dealers who made markets in Abbey National (ANL), a FTSE stock in 1991. Each graph records the dealer's inventory (in shares) versus calendar time. During 1991, ANL had an average price of £2.70 and a NMS of 50,000 shares. (Thus, dealers registered to quote prices for ANL were required to take trades of up to 50,000 shares, or £135,000, at their posted quotes.) To see the connection between inventories and interdealer trades, we have superimposed circles and stars on the inventory series. Circles indicate a dealer's inventory just prior to an IDB trade; stars indicate the inventory just prior to an IMM trade. We have not marked customer or broker trades, which account for the remaining variation in inventories.

Panel A graphs the inventory for a dealer from a small firm who tightly manages his position. Most trades are below the NMS size of 50,000 shares. Rarely does this dealer accept trades above 50,000 shares. When the dealer's position exceeds roughly $\pm 10,000$ shares, he often resorts to interdealer trades to move back toward zero inventory. Panel B shows the inventory of a dealer from a medium-sized firm with limited customer order flow. He appears to use interdealer trades, particularly IDB trades, to take positions. This dealer takes positions of up to 700,000 shares (£1,890,000!) in either direction during the course of the year. He usually unwinds positions within four to six weeks.

The dealer in Panel C also frequently uses interdealer trades and has substantial customer order flow. In contrast to dealer B, he does not take positions. Instead, he uses interdealer trades to unwind large customer trades. From the figure we also can see that when the dealer has a large position to unwind, he splits his interdealer trades. (One can see this as a sequence of circles or stars.) The dealer in Panel D trades for one of the most active London firms. He takes large positions, often in excess of 1 million shares. He also actively uses interdealer trades to accumulate positions, which are then reversed. Like dealer C, he attracts a substantial amount of customer order flow. It thus appears he uses interdealer trades to take both long-term positions in advance of customer trades and to undo the inventory consequences of customer trades.

B. Definition of Inventory Cycles

What should we make of the inventory management practices of dealers A, B, C, and D? Inventory imbalances for dealers A and B seem to have been primarily caused by block trades (thus the vertical spikes), but dealer C is

¹² Thus, if a dealer has an average end-of-day inventory of 25,000 shares during 1991, the inventory series we work with is the dealer's inventory (as calculated by us) minus 25,000 shares. The purpose of this adjustment is to remove an estimate of the unknown (to us) inventory the dealer had at the start of the year.




Panel A**Panel B**

Figure 1. 1991 Dealer Inventory in Abbey National plc. The figure displays the actual 1991 intraday inventory series for four of the fourteen dealers registered as market makers in Abbey National. We label the dealers A, B, C, and D for confidentiality reasons. Each trade for the dealer's own account is cumulated starting from zero to create an inventory series. Inventories are then demeaned by subtracting each dealer's 1991 average daily closing inventory from each observation. The resulting inventory series is represented by a solid line. The inventory prior to an IDB trade is marked with a circle, and the inventory prior to an IMM trade is marked with a star.

Panel C



Panel D




Figure 1. Continued.

apparently actively taking positions unrelated to block trades—to a large extent using interdealer trades (circles and stars). The inventory series for dealer D displays a mixture of these two types of inventory imbalances. The challenge is to develop a methodology that enables us to account for this heterogeneity in trading strategies.

The empirical literature on inventory management summarizes the behavior of inventory series by calculating the mean reversion of daily inventories (see, e.g., Stoll (1976), Hasbrouck and Sofianos (1993), and Madhavan and Smidt (1993)). Using 1991–1992 London data, Hansch et al. (1998) estimate that dealers on average unwind half their normalized inventories in 2.5 trading days. They note that this is substantially faster than previous studies using NYSE specialist data have found.¹³ We see two main drawbacks to focusing on dealers' daily inventories. First, as Figure 1 shows, inventories vary considerably, both across dealers *and* across time. Second, inventories vary considerably within the trading day.

Departing from tradition, we base our analysis of dealers' inventories on the idea of an inventory cycle. Roughly, a cycle is a sequence of trades in which a dealer's inventory departs from a "desired" long-run range, reaches an "extreme," and then returns to the desired range. As many dealers tell it, most cycles are the result of one order or a sequence of orders that attract the dealer's attention because their inventory reaches an unusual level (either short or long). At this point, the dealer seeks to reduce the imbalance by relying on a cost-effective combination of inventory management tools: soliciting customer orders (by phone or using quotes); using IDB screens; and calling competing dealers (IMM). The inventory series of dealers A and B in Figure 1 illustrate typical imbalances where one or a few large customer or broker (block) trades cause an imbalance. The two key practical issues we face when defining an inventory cycle are: what do we mean by "extreme" and what do we mean by "desired" inventories?

A dealer's inventory management policy in principle depends on the risk aversion of the dealer and firm, and on the firm's capital resources. Since we do not have this information, we need a universal classification scheme. We assume that zero inventory is the "desired" long-run inventory level of all dealers.¹⁴ We define an "extreme" long position as a sequence of trades that cause the dealer's inventory to pass three critical levels. These levels are labeled 2, 3, and 4 in Figure 2. Point 3 corresponds to an inventory level equal to three times the stock's NMS. (Recall that one NMS is the mandatory volume that a dealer must guarantee at his posted prices.) We take this as a necessary condition for a long cycle because, assuming dealers would grant counterparties at most one NMS at the public market quotes, it would take the dealer at least three interdealer trades to return to zero inventory. Three NMS is also the minimum customer trade size for which trade publication is automatically delayed by 90 minutes during our sample period.

¹³ For example, Madhavan and Smidt (1993) estimate the average inventory half-life to be almost 50 days unless shifts in desired inventory levels are allowed, which brings the average down to 7.3 days.

¹⁴ When we say zero is the "desired" inventory level, we mean that the dealer expects the inventory to fluctuate about zero, but does not make long-run bets on a stock. To the extent that dealers do take short-run speculative positions, our ability to reject the null will be reduced.




Figure 2. The cycle algorithm. The figure displays a hypothetical dealer inventory series. Each trade for the dealer's own account is represented by the solid line. Our algorithm assumes that the dealer starts managing his inventory after it exceeds ± 3 NMS. For example, when the dealer receives a large customer sale (vertical segment) that takes his inventory from 1 NMS (point 2) to just above 3 NMS (point 3), he starts working on reducing his inventory back to one NMS (point 4). The dealer subsequently faces a sequence of customer buy orders, which reduce his inventory past -1 NMS (point 6) to -3 NMS (point 7). He then works this short inventory position back to the balanced region (point 8).

Points 2 and 4 identify the boundary of the desired inventory range, and thereby define the beginning and end of an inventory cycle. Figure 2 displays two such cycles. The shaded regions constitute the set of trades that we analyze below. The first (long) cycle is caused by one large customer trade, the second (short) cycle illustrates an imbalance resulting from a sequence of customer trades. We choose to start and end inventory cycles at a one NMS band about zero inventory because one NMS is the amount the Exchange requires a dealer to accept if the dealer is posting at the touch. Further, within this band, a dealer could instantaneously return to zero inventory by hitting a competing dealer's quotes.¹⁵ One final note, many trades are not part of a cycle. For example, we do not analyze trades between points 1 and 2 because the dealer's inventory does not reach three NMS.

An implicit assumption underlying our description of an inventory cycle is that the dealer takes one or a few large block trades in one direction and *then* uses interdealer trades to reverse the imbalance. Most of our large inventory imbalances follow this sequencing. We refer to this type of trade pattern as an "unanticipated" imbalance. Some block trades are likely to have been anticipated by the dealer, however, because London dealers ex-

¹⁵ The NMS boundaries differ by security and dealer. During 1991 there were two Scottish dealers who were allowed to quote one-half a stock's NMS. We therefore adjust their desired trading range to be one-half the normal NMS. In future work, we would like to allow the cycle boundaries to vary by dealer.




Figure 3. Unanticipated and anticipated cycles. The figure displays a hypothetical dealer inventory series. Each trade for the dealer's own account is represented by the solid line. The first inventory spike starts with a large customer sell order (vertical line) followed by a sequence of interdealer, customer, and broker trades that reduce the inventory back to the ± 1 NMS region. We call this type of cycle an unanticipated customer trade-related cycle because the large customer order was likely a surprise. By contrast, an anticipated cycle has a large customer buy order (vertical line) that is preceded by a sequence of interdealer, customer, and broker trades. In this case, the dealer is likely to have known about the arrival of the large customer order in advance, and thus "worked" the order.

ecute "protected" trades. Protected trades involve a dealer offering potential price improvement to a customer requesting a large trade, provided that the dealer gets some time to "work" the order. Indeed, the Exchange estimates that roughly 30 percent of large customer trades are "worked" orders.¹⁶

Figure 3 illustrates what we have in mind. The first (in time) long inventory cycle consists of one block trade, followed by a sequence of IDB trades (the circles). This is the most common interdealer trading pattern and we call it an *unanticipated cycle*. The second long inventory cycle is different; it consists of a sequence of IDB trades (buys) followed by one block trade (a customer or broker buy). We refer to this second cycle as an *anticipated cycle*. An anticipated cycle is one where the customer is likely to have called the dealer and stated his trading interest. Instead of requesting immediate execution, however, he allows the dealer to delay "printing" the large order until the dealer has positioned the large order.

¹⁶ The dealer can work the order until the end of the day and then the trade must be printed and revealed to the market. (See the LSE consultation document, *New Electronic Trading Service*.) These "worked" orders have become so important that they now receive special treatment in London's new trading system, the Stock Exchange Electronic Trading Service (SETS); see *Rules for the Stock Exchange Electronic Trading Service*, Stock Exchange Notice N4/97.

IV. Empirical Evidence on Risk Sharing

In this section, we use 1991 data from the London Stock Exchange to test our five hypotheses. Using the algorithm described in Appendix B, we first identify all inventory cycles for all dealers in each of our 25 FTSE and 20 non-FTSE stocks. We then identify inventory cycles caused by block trades. Our algorithm dictates that at least two-thirds of the side of the cycle that contains the customer (broker) trades that produced the imbalance be customer (broker) volume. We do not constrain the composition of trading volume on the side that reverses the imbalance.

Table II reports the characteristics of our sample cycles by FTSE and non-FTSE stocks.¹⁷ Our algorithm identifies 10,428 cycles for FTSE stocks and 6,140 cycles for non-FTSE stocks. There are 17 dealing firms active for FTSE stocks and 16 dealing firms active in our non-FTSE sample. Thus, on average, FTSE (non-FTSE) dealing firms experience 613 (384) cycles in total, or 25 (20) cycles per stock, during 1991. The table reveals that 83 to 84 percent of all cycles are caused by customer trades. An additional 1 to 2 percent are due to broker trades, and 4 and 9 percent represent crosses in which the dealer pairs two large offsetting trades. These crosses resemble large trades negotiated in the upstairs market of the NYSE. The remaining cycles (6 percent for non-FTSE stocks and 13 percent for FTSE stocks) represent a residual category. They are cycles that we cannot classify confidently as primarily related to customer or broker trades. We thus label them ambiguous cycles. These cycles could reflect customer or broker trade-related cycles, mistimed cycles, or cycles in which dealers used interdealer trades to take speculative positions. Most dealers have at least a few such ambiguous cycles. Thus, it is not simply a matter of some dealers specializing in taking positions unrelated to block trades, nor does it appear our algorithm misclassifies cycles for any particular dealer.

Table II also reports on how many cycles are anticipated, the median duration of each type of cycle, the time elapsed between the start of a cycle and its peak, and the time elapsed between the peak and the finish of a cycle.¹⁸ Roughly one-third of the cycles caused by customer trades are anticipated. This is strikingly similar to Exchange claims that roughly 30 percent of large orders are “worked” (Oliver Wyman & Co. (1996)). More than half of the customer trade-related cycles in FTSE stocks last less than one day, and thus would not be captured in daily inventory data. Cycles for non-FTSE stocks last longer, one day if anticipated and two if not. Broker trade-related cycles are similar, although they are considerably shorter for non-FTSE stocks. By comparison, the cycles we cannot classify as customer or broker trade-related are considerably longer. The ambiguous FTSE and non-FTSE cycles average six and five days to complete, respectively. As one might expect if

¹⁷ We use medians because our algorithm that estimates cycle inventories is subject to error. The median reduces the possible influence of data outliers or classification errors.

¹⁸ The units are 24-hour days. Minimum trading hours in London are from 8:30 a.m. to 4:30 p.m. Our algorithm captures all trades, whether they take place before, during, or after hours.

Table II
Characteristics of Inventory Cycles: Sample Medians

This table reports statistics for the inventory cycles of the twenty-five FTSE and twenty non-FTSE stocks in our sample. We define an inventory cycle as a period that starts and ends when a dealer's inventory crosses one NMS (normal market size) in absolute value. We require that during the cycle the dealer's inventory must exceed three NMS (in absolute value). A security's NMS equals 2.5 percent of an average day's trading volume. Customer trades are trades where at least one counterparty is someone other than a London dealer or broker. Broker refers to proprietary trades by Exchange brokers who are not dealers in the stock. For a cycle to be well defined, we require that at least 67 percent of the side creating the imbalance consist of customer (broker) trades. Unanticipated cycles occur when a dealer unwinds one or a few large customer (broker) orders. Anticipated cycles occur when a dealer trades ahead of one or several large customer (broker) orders. Ambiguous cycles are cycles that we cannot classify as anticipated or unanticipated. We divide ambiguous cycles into cycles that would be classified as customer related if we classified them as having more than 50 percent customer volume on the side causing the cycle (as opposed to 67 percent). Crosses occur when two nearly identical customer trades are matched by a dealer in a short time. We measure cycle lengths in 24-hour days. The peak time of a cycle is the time stamp of the highest (lowest) inventory for a long (short) cycle. The overall fraction of customer volume equals customer volume as a fraction of the total number of shares bought and sold by the dealer during the cycle. Peak volume is the maximum number of shares the dealer unwound or built during a cycle. Peak volume is different from the size of the cycle, which is the cumulative number of shares the dealer unwound (or built) relative to a position of ± 1 NMS. We identify the side generating the cycle based on an algorithm for classifying unanticipated and anticipated cycles. For ambiguous cycles and crosses, we designate dealer purchases (sells) as trades on the side generating long (short) positions.

Cycle type	Total Number of Cycles	Cycle Length			Overall		Side Generating Cycle		
		Percent of Cycles	Finish- Start (days)	Peak- Start (days)	Finish- Peak (days)	Fraction of Customer Volume	Absolute Peak Volume (NMS)	Fraction of Customer/ Broker Trades	Number of Trades
FTSE stocks	10,428								
Unanticipated customer	5,598	53.7	1.00	0.02	0.75	0.60	5.6	1.00	3
Anticipated customer	3,006	28.8	0.85	0.13	0.04	0.59	5.6	1.00	2
Unanticipated broker	80	0.8	0.86	0.04	0.73	0.12	4.4	0.94	2
Anticipated broker	41	0.4	0.86	0.12	0.05	0.08	4.6	0.97	1
Crosses	382	3.7	0.00	0.00	0.00	1.00	5.0	1.00	1
Ambiguous	1,321	12.7	5.89	1.74	2.78	0.36	5.6	0.35	9
Potential unanticipated	381	3.7	7.00	2.87	3.01	0.46	6.9	0.61	13
Potential anticipated	282	2.7	7.04	1.97	3.19	0.47	6.7	0.60	14
Residual	658	6.3	4.28	1.01	2.03	0.22	4.8	0.23	6
Non-FTSE stocks	6,140								
Unanticipated customer	3,405	55.5	1.99	0.04	0.95	0.69	6.6	1.00	2
Anticipated customer	1,741	28.4	0.95	0.15	0.06	0.71	6.9	1.00	1
Unanticipated broker	57	0.9	0.20	0.00	0.13	0.02	5.0	1.00	1
Anticipated broker	49	0.8	0.05	0.00	0.01	0.01	4.3	1.00	1
Crosses	573	9.3	0.00	0.00	0.00	1.00	7.6	1.00	1
Ambiguous	315	5.1	5.20	1.80	2.15	0.34	5.9	0.32	5
Potential unanticipated	96	1.7	6.84	2.09	2.99	0.43	7.1	0.59	5.5
Potential anticipated	54	0.9	8.89	2.47	4.31	0.48	7.1	0.61	7
Residual	165	2.7	4.01	1.05	1.77	0.23	5.3	0.25	4

crossing trades are prearranged, these trades typically take little time to complete. Moreover, the data suggest that it typically takes three to five hours for London dealers to work a large customer order.

The right-hand side of Table II includes information on the median fraction of total customer (or broker) volume traded (buys plus sells) and the peak cycle volume. For customer trade-related cycles, customer trades account for 60 to 70 percent of the volume. This is slightly lower than the overall sample proportion reported in Table I. By contrast, ambiguous cycles have considerably less customer volume. That is, they contain a disproportionate number of interdealer trades. Cycles caused by broker trades have even less customer volume (1 to 12 percent). By definition, crosses have 100 percent customer volume. The median peak inventory ranges from four to seven NMS for both FTSE and non-FTSE stocks. In the case of ANL, a position of five NMS corresponds to 250,000 shares, or roughly £675,000.

The final two columns in Table II report information on the trades that we classify as causing the cycle. For customer (broker) trade-related cycles we require two-thirds of the volume on the side generating the imbalance to be customer (broker) trades. According to Table II, the median customer trade-related cycle consists of two to three customer orders. These customer orders tend to be large. In results not reported here, we find that roughly two-thirds of our customer cycles involved at least one large customer trade of three NMS or more. We choose not to analyze these cycles separately because we find virtually no qualitative differences in interdealer trade patterns between cycles caused by one large trade and those caused by many customer trades. From Table II one can also see that broker trade-related cycles tend to consist of one or a few large broker trades. The ambiguous cycles that we do not classify as customer or broker trade-related have much less customer and broker volume, and many more trades. If we relax the requirement that two-thirds of one side of the cycle be customer or broker volume, and make it half of one side, we classify more cycles (labeled *Potential anticipated* and *Potential unanticipated* in Table II). These “potential” cycles tend to last longer and have more trades, but are otherwise qualitatively similar to customer and broker trade-related cycles. Because we are less comfortable applying our anticipated and unanticipated classification to these cycles, we only briefly summarize results for these cycles.

A. Does Interdealer Trading Increase during Inventory Imbalances?

Our first hypothesis asserts that IDB trading should increase when the dealer's inventory increases. Table I suggests that, on average, interdealer trade accounts for 24 percent of all trades. Hence, we take as our null that we should see approximately 24 percent of trades being interdealer trades during inventory cycles. To summarize the extent of interdealer trading during inventory cycles, we calculate in Table III how much of a dealer's posi-

Table III
Trading Patterns in Inventory Cycles

This table reports statistics for the inventory cycles of the twenty-five FTSE and twenty non-FTSE stocks in our sample. We define an inventory cycle as a period that starts and ends when a dealer's inventory crosses one NMS (normal market size) in absolute value. We require that during the cycle the dealer's inventory must exceed three NMS (in absolute value). A security's NMS equals 2.5 percent of an average day's trading volume. Unanticipated cycles occur when a dealer unwinds one or a few large customer (broker) orders. Anticipated cycles occur when a dealer trades ahead of one or several large customer (broker) orders. We exclude ambiguous cycles from this table (see Table II). We define the dealer's position to be the sum of all shares bought (or sold), starting from an inventory position of ± 1 NMS. IDB refers to trades consummated on one of the four interdealer broker systems. IMM refers to direct trades between dealers registered to quote prices for a security. Broker refers to trades by Exchange brokers who are not also dealers. Broker trades include crosses or trades between a broker and a dealer.

Cycle type	Number of Cycles	Average Fraction of Position Created by				Average Fraction of Position Reversed by			
		IDB Trades (%)	IMM Trades (%)	Broker Trades (%)	Customer Trades (%)	IDB Trades (%)	IMM Trades (%)	Broker Trades (%)	Customer Trades (%)
FTSE stocks									
Unanticipated customer	5,598	1.9	2.5	0.5	95.1	32.8	29.2	2.5	35.5
Anticipated customer	3,006	2.9	2.3	0.4	94.4	30.9	31.0	3.8	34.3
Unanticipated broker	80	1.3	1.4	90.3	7.0	41.5	34.8	2.5	21.1
Anticipated broker	41	0.6	2.3	90.0	7.1	43.8	30.2	6.0	20.0
Non-FTSE stocks									
Unanticipated customer	3,405	1.0	1.6	0.3	97.1	26.3	23.9	3.7	46.1
Anticipated customer	1,741	1.2	1.2	0.3	97.4	25.8	21.3	6.0	46.9
Unanticipated broker	57	0.0	1.7	94.3	4.0	37.9	49.5	2.3	10.3
Anticipated broker	49	0.4	0.8	95.6	3.2	48.9	34.3	2.9	13.9

tion was created and reversed, respectively, by customer and interdealer trades. The rows of the table distinguish between cycles resulting from trades that are anticipated versus unanticipated. We also group the results by FTSE and non-FTSE stocks. To illustrate our calculations, consider a dealer who is long because of a large ten NMS unanticipated customer sell order. Assuming that the customer order is the only sale to the dealer during the inventory cycle (that is, the dealer started with no inventory), the “size” of the position the dealer must reverse is nine NMS (the ten NMS trade less the one NMS band). The primary ways the dealer can reverse his long position are sales to customers, sales to brokers, and sales to competing dealers (the second set of columns of Table III).

First, consider inventory cycles caused by large customer or broker trades (unanticipated and anticipated). It is clear from the first set of columns that these cycles are primarily caused by customer and broker trades, respectively. Indeed, 94 to 97 (90 to 96) percent of the volume generating a position is customer (broker) trades (cf. Table II). Moreover, Table III shows that during customer trade-related cycles in FTSE stocks, approximately 65 percent of the dealer’s inventory is reversed by interdealer trades. This is true whether the customer trade(s) is (are) anticipated or unanticipated. Though interdealer trades in non-FTSE stocks account for a somewhat smaller percentage, they nevertheless are used more frequently than customer trades to reverse extreme positions. Cycles caused by broker trades have an even larger preponderance of interdealer trades during the reversal—80 percent for FTSE stocks and roughly 87 percent for non-FTSE stocks. If we distinguish between the different types of interdealer trades, we see that both IDB and IMM trades roughly triple during customer trade-related cycles. By contrast, broker trades are used infrequently. For broker trade-related cycles, IDB volume and IMM volume similarly increase. Dealers’ preferences between IDB and IMM trades appear to differ somewhat for FTSE and non-FTSE stocks. In part, IDB trades tend to be less common in the smaller stocks because the IDBs have less liquidity in smaller stocks.

Finally, consider what would happen if we do not distinguish between anticipated and unanticipated inventory imbalances. For anticipated cycles, the entries in the columns labeled “Average Fraction of Position Created” would then be replaced by what currently is in the columns labeled “Average Fraction of Position Reversed,” and vice versa. For example, for anticipated customer trade-related cycles in FTSE stocks, we would mistakenly report that IDB, IMM, and customer trades each created roughly one-third of the dealer’s position and that customer trades reversed 94.4 percent of the dealer’s imbalance. Since nearly 30 percent of our cycles appear to be the result of worked orders, if we inadvertently pooled anticipated with unanticipated cycles, we would therefore drastically underestimate the importance of interdealer trading.¹⁹

¹⁹ A rough calculation of the downward magnitude for customer-induced cycles is $0.3 \times 5\% + 0.7 \times 60\% = 43.5\%$, versus the roughly 60 percent of interdealer trades we observe in Table III.

A concern one might have about the fractions reported in Table III is that they might obscure considerable heterogeneity in interdealer trading during extreme inventory cycles. Some cycles have no interdealer trading and others consist entirely of interdealer trades. Part of this heterogeneity may be due to the way we measure cycles. Specifically, our three NMS rule used to define cycles may include cycles that are not “extreme” in the sense of Hypothesis 1. If this were true, we might expect to find that the extent of interdealer trading would increase with the amplitude of the cycle (the amount of trading required to return to one NMS). One also might expect that the definition of “extreme” would be sensitive to the inventory positions of other dealers. That is, inventory risk, and hence interdealer trades, might increase when the average exposure of all dealers increases.

Because inventory cycles differ across dealers and through time, we estimate a multivariate statistical model relating the fraction of cycle volume reversed by interdealer trades to cycle, stock, and dealer characteristics. These models only analyze the cycles that are caused by large customer trades. The broker trade sample is too small to lend itself to this kind of analysis. Because the dependent variable is truncated at 0 and 1, we use a doubly truncated Tobit model. Table IV summarizes several of these Tobit regressions. As is common with limited dependent variable models, it makes more sense to report estimates of the expected change in the dependent variable (the fraction of cycle trade that was interdealer trade) from a one unit change in each regressor (see Greene (1997).) This is the approach taken in Table IV, Panel A. To further aid in the interpretation of the coefficients, we supplement the table with predictions of the fractions of interdealer trade that we would observe for hypothetical cycles in Panel B.

We include a range of regressors in the regressions in an effort to isolate interesting intercycle variations in the fraction of interdealer trade. The variable UNANT is a zero-one dummy variable equal to one when the cycle is unanticipated; ANT is a zero-one dummy variable equal to one when the cycle is anticipated; FTSE is a zero-one dummy variable equal to one when the stock is a FTSE issue; NONFTSE is a zero-one dummy variable equal to one when the stock is a non-FTSE issue; SIZE is the peak cycle inventory that the dealer positioned; TOTAL is the sum of all dealers' inventories when the cycle first reaches three NMS in the case of anticipated cycles or crosses three NMS the second time in the case of unanticipated cycles; and ID SELLER and ID BUYER are zero-one dummy variables equal to one when the dealer is buying or selling, respectively. We include SIZE because we expect it to be difficult and time consuming to unwind a very large customer order. We include TOTAL because we believe that the aggregate inventory of all dealers should influence the extent of interdealer trading. We interact TOTAL with ID SELLER and ID BUYER for the following reason. Consider a dealer who anticipates having a long position. This dealer would like to sell. We expect, however, that the dealer is less likely to use interdealer trades during the downward part of the cycle

Table IV
**Intercycle Variation in Interdealer Trade:
 Estimated Marginal Effects of Cycle Size and Positions**

Panel A reports results from a two-limit Tobit model that explains the fraction of interdealer trade reversing a customer trade-related inventory imbalance. Panel B uses the two-limit Tobit coefficient estimates to predict the percentage of inventories that will be reversed by interdealer trades. The estimates in Panel A represent the increase in the fraction of interdealer trade associated with a one-unit increase in the corresponding independent variable. These estimates take into account the truncation of the dependent variable at zero and one (see Greene (1997)). The intercept estimate is the original intercept from the Tobit model. The numbers in parentheses are the estimated (asymptotic) *t*-statistics. We have classified inventory cycles into *anticipated* cycles in which interdealer trading precedes customer trades and *unanticipated* cycles in which interdealer trading follows customer trades. FTSE is a zero-one dummy variable equal to one for FTSE stocks. NONFTSE is a zero-one dummy variable equal to the complement of FTSE. UNANT is a zero-one dummy variable equal to one if the dealer is positioning an unanticipated customer trade using interdealer trades. ANT is a zero-one dummy variable equal to one if the dealer is positioning an anticipated customer trade using interdealer trades. ID BUYER is a zero-one dummy variable equal to one if the dealer is positioning by buying shares. ID SELLER is a zero-one dummy variable equal to one if the dealer is positioning by selling shares. NMS is a security's normal market size; one unit equals 2.5 percent of an average day's trading volume. SIZE is the cycle inventory position (in NMS units) that the dealer is unwinding using interdealer trades. TOTAL is the sum of all dealers' inventories at the point in the cycle when a dealer is at three NMS and trading to unwind his position. IDB refers to trades on one of the four interdealer broker systems. IMM refers to direct trades between dealers registered to quote prices for a security.

Panel A: Estimated Marginal Effects of Cycle Size and Positions

	Dependent Variable Fraction of Inventory Position Reversed by		
	All Interdealer Trades	IDB Trades	IMM Trades
CONSTANT	0.5877 (47.73)	0.0751 (4.62)	-0.0205 (-1.21)
UNANT	0.0229 (2.22)	0.0206 (1.93)	0.0545 (5.12)
FTSE	0.1580 (15.32)	0.1236 (10.60)	0.2069 (17.31)
FTSE*UNANT	-0.0354 (-2.72)	0.0222 (1.69)	-0.0879 (-6.59)
FTSE*UNANT*SIZE	-0.0033 (-18.53)	-0.0004 (-2.11)	-0.0000 (-0.24)
NONFTSE*UNANT*SIZE	-0.0024 (-13.95)	-0.0004 (-2.62)	-0.0009 (-5.75)
FTSE*ANT*SIZE	-0.0027 (-11.86)	-0.0003 (-1.61)	-0.0006 (-2.63)
NONFTSE*ANT*SIZE	-0.0021 (-10.52)	-0.0010 (-4.92)	-0.0003 (-1.73)
FTSE*ID SELLER*TOTAL	-0.0004 (-4.04)	-0.0009 (-7.69)	0.0003 (2.56)
NONFTSE*ID SELLER*TOTAL	-0.0009 (-5.74)	-0.0012 (-6.31)	-0.0004 (-2.38)
FTSE*ID BUYER*TOTAL	0.0009 (7.85)	0.0009 (7.18)	0.0000 (0.39)
NONFTSE*ID BUYER*TOTAL	0.0012 (7.55)	0.0016 (9.93)	0.0005 (3.01)

Table IV—Continued

Panel B: Predicted Percentages of Inventories Reversed by Interdealer Trades during Inventory Cycles		Predicted Percentage of Inventory Position Reversed Through		
		All Interdealer Trades	IDB Trades	IMM Trades
Sample means	Full sample	60.8	30.4	27.4
	FTSE sample	65.4	32.1	30.0
	Non-FTSE sample	55.1	26.5	23.4
Estimated means	FTSE stocks	67.2	33.2	31.4
holding trade	<i>Anticipated</i> 10 NMS inventory;			
characteristics	dealer is <i>buying</i> ; all dealers			
constant	long a total of 30 NMS			
	<i>Unanticipated</i> 10 NMS inventory;	66.0	35.7	29.9
	dealer is <i>buying</i> ; all dealers			
	long a total of 30 NMS			
	<i>Unanticipated</i> 10 NMS inventory;	63.1	32.6	30.2
	dealer is <i>selling</i> ; all dealers			
	long a total of 30 NMS			
	Non-FTSE stocks			
	<i>Anticipated</i> 10 NMS inventory;	57.4	27.4	21.9
	dealer is <i>buying</i> ; all dealers			
	long a total of 30 NMS			
	<i>Unanticipated</i> 10 NMS inventory;	58.8	28.8	24.1
	dealer is <i>buying</i> ; all dealers			
	long a total of 30 NMS			
	<i>Unanticipated</i> 10 NMS inventory;	54.2	24.5	22.9
	dealer is <i>selling</i> ; all dealers			
	long a total of 30 NMS			

if the aggregate position of all dealers is long ($TOTAL > 0$) than if the aggregate position of all dealers is short ($TOTAL < 0$).²⁰ Accordingly, larger values of $TOTAL$ should reduce interdealer trading when the dealer is selling. The opposite should be true when the dealer is buying.

The coefficient estimates for $UNANT$ and $FTSE$ in Table IV largely confirm the relations among the conditional means in Table III. To understand the interaction effects with $TOTAL$, we subsequently report the fractions of dealers' inventories that would be reversed by interdealer trades under different trading scenarios. The scenarios show that $FTSE$ dealers more frequently unwind their positions with interdealer trades than non- $FTSE$ dealers. We also find that dealers are just as apt to use interdealer trades when they work an order (anticipated inventory imbalance) as when they position an unanticipated inventory imbalance. Though Table III suggests dealers are equally likely to reduce inventories with IMM and IDB trades, Table IV suggests that dealers reverse a greater fraction of imbalances using IDB trades. This finding confirms Hypothesis 1. Compared to other cycles, we see that a $FTSE$ dealer who buys to reverse a short position of ten NMS , and who faces competitors who on average are long, will use interdealer trades to reduce two-thirds of his peak position (6.6 NMS). This fraction declines about three percent if his competitors have similar inventory positions—that is, they too are short the stock. Thus, unlike Manaster and Mann (1996), we find that aggregate inventory affects dealer trading in the direction predicted by inventory or risk-sharing models. The results in Table IV also suggest that as the size of a dealer's inventory imbalance increases, the percentage reversed by interdealer trades declines slightly. Most of this decline comes about because of a fall in the percentage of IMM trades. Relative to the overall variation in interdealer trading, the impact of position size on the fraction of interdealer trade is small. For instance, an increase in a dealer's position by ten NMS (roughly 1.5 times the sample average) would decrease the percentage of the inventory reduced by interdealer trades by about 2.5 percent.

We find the results in Tables III and IV striking. We reject the null hypothesis that interdealer trade is independent of dealer inventory levels during cycles caused by customer or broker trades. Indeed, interdealer trades are on average used more often than customer trades to reduce extreme inventories resulting from one or several customer or broker trades. Hansch et al. (1998) also provide support for this hypothesis based on a different methodology. We also find evidence that dealers tend to prefer IDB trades over IMM trades when IDB systems are likely to be liquid, presumably because of their relatively lower execution costs. Interdealer trading activity also is much more common in $FTSE$ stocks even when we control for the overall size of the imbalance. One explanation for this difference is that $FTSE$ stocks have more dealers. A greater number of dealers can increase

²⁰ This argument overlooks the fact that the aggregate market position could reflect a mix of dealers working orders and positioning unanticipated customer trades.

dealers' opportunities for risk sharing. While IDB, IMM, and customer trades are roughly equally important for building positions in our unclassified (residual) category, IDB trades are relatively disfavored in reversing positions. A possible explanation is that dealers prefer a rapid elimination of at least some of these positions (potentially because of anticipated public release of information on which the position was built).

Finally, in other regressions not reported here, we test whether particular dealers tend to use interdealer trades more often than others. Though we do not report specific results for individual dealers or firms, we do find significant differences across firms. In particular, it appears that the largest six firms (often referred to in London as the "Leaders") are substantially less likely to use interdealer trades. For example, on average they unwind 15 percent (by value) more of their positions with customer trades.²¹ This finding might suggest that the "Leaders" have different risk tolerances (or capital) than other dealers, but it could also reflect the fact that these firms have greater access to preferenced order flow. Taken together, these findings support Hypothesis 1: interdealer trading increases when a dealer's position is extreme.

B. Do Dealers Use Interdealer Trades to Reduce Positions?

The previous subsection shows that interdealer trade increases when dealers have (or build) large positions. Though this evidence suggests that interdealer trade may facilitate inventory management and risk sharing, we have yet to show that the direction of interdealer trades is consistent with inventory management. If interdealer trades had nothing to do with a dealer's inventory position, then we would expect that no matter what a dealer's position, 50 percent of the dealer's interdealer trades would be buys and 50 percent would be sales (the population proportions). If Hypothesis 2 is correct, however, then we would expect to see interdealer trades follow particular patterns. These patterns depend both on the direction of the dealer's position (long or short) and whether the position is anticipated or unanticipated.

To make these four possibilities clear, let us first consider what happens when the dealer's position is unanticipated. If the dealer receives one or more customer sales and thus has a long position, we would expect the dealer to use interdealer sales to reduce his inventory risk. Thus, we should not see any interdealer buys during such a cycle.²² Conversely, if the dealer acquires an unanticipated short position through customer buys, the inventory hypothesis would predict that the dealer would use interdealer buys to increase his inventory and the cycle should not include any interdealer sells. It is important to note that these predictions about the direction of interdealer trade hinge critically on the dealer's position being *unanticipated*.

²¹ Each of the "Leaders" makes markets in over 1,000 stocks. They also typically average 20 percent of the customer business in each stock in which they are active.

²² Since a dealer can end up executing an IMM trade because another dealer hits his quotes, we may well see some IMM trades in the "wrong" direction.

Table V
Interdealer Trading to Unwind Positions

This table reports results for all interdealer trades in which at least one dealer is in an inventory cycle. If both dealers are in an inventory cycle, we count the trade twice. We exclude cycles resulting from broker trades. We define *anticipated* inventory cycles as cycles where interdealer trading precedes customer trades and *unanticipated* cycles as cycles where interdealer trading follows customer trades. IDB refers to trades on one of the four interdealer broker systems. IMM refers to direct trades between dealers registered to quote prices for a security.

Cycle Type	FTSE Stocks		Non-FTSE Stocks	
	Number of Trades	Percent of Trades Unwinding	Number of Trades	Percent of Trades Unwinding
Unanticipated customer				
IMM trades	10,527	79.3	4,149	82.1
IDB trades	14,668	82.0	4,213	86.5
Anticipated customer				
IMM trades	5,391	81.3	1,974	84.8
IDB trades	7,271	76.4	1,874	81.0

What happens if the dealer instead *anticipates* having a position? If the broker-dealer anticipates executing customer orders (protected trades) that will result in a position, then the dealer has the option to build an offsetting position (by working the order) prior to reporting the customer trade. In this case, we expect interdealer buys to *precede* anticipated customer buys and interdealer sales to *precede* anticipated customer sales. We also would expect there not to be interdealer sells (buys) preceding anticipated customer buy (sell) orders. Thus, whether or not the dealer anticipates the trade has an important impact on the direction of interdealer trades that would reduce inventory risk.

Table V reports the fraction of all interdealer trades taking place during cycles that are in the direction predicted by our risk-sharing hypothesis.²³ Again, if interdealer trades had nothing to do with dealer inventories, we would expect a 50-50 split between buys and sells. The table reports results for cycles caused by anticipated and unanticipated customer trades. For example, 79.3 percent of all IMM trades in unanticipated FTSE cycles were in the right direction. That is, they were sells for dealers in long cycles and buys for dealers in short cycles. Similarly, 82.0 percent of all IDB trades were in the direction predicted by inventory risk sharing. We obtain similar results for anticipated cycles. One difference is that for anticipated cycles, we find IMM trades in the predicted direction more often than IDB trades.

²³ Interdealer trades take place outside of cycles for at least three reasons. First, we exclude interdealer trades in which dealers have yet to begin a cycle at the start of 1991 or are in the midst of a cycle at the end of 1991. Second, our algorithm may misclassify cycles. Third, in the case of IMM trades, Exchange rules may force a dealer to trade with another.

Given the large number of cycles underlying our percentages, the fractions in Table V easily reject the 50-50 null hypothesis of no inventory motive for interdealer trades during cycles caused by customer trades.²⁴ Across both FTSE and smaller stocks, we find that roughly 80 percent of all interdealer trades during customer trade-related cycles reverse dealers' positions. The results for cycles caused by broker trades (not reported) are very similar. While we take this as strong evidence for the hypothesis that interdealer trades help reduce dealer inventory risk (Hypothesis 2), we also can reject the hypothesis that all interdealer trades during these cycles are inventory related. We believe, however, that measurement and classification errors are likely to bias our calculations toward the 50-50 null hypothesis. For instance, if our anticipated/unanticipated cycle classification scheme is imperfect, we will misclassify trade directions. The fact that a dealer can end up in an IMM trade inadvertently when a competitor hits his quotes also works against us. Moreover, if dealers post orders on IDB screens to try to influence competitors, and those orders are hit, we are unlikely to find 100 percent of IDB trades in the direction predicted by a risk-sharing hypothesis.

C. Do Dealers with the Most Extreme Inventories Trade with Each Other?

In most complete information risk-sharing models the degree of risk sharing critically depends on participants' degrees of risk aversion and the distribution of risky endowments (i.e., inventories). One could, for example, imagine two dealers with identical inventory positions trading with one another to reduce risk. This could happen if one dealer is more risk averse than the other. If we assume, as is done in most of the literature, that dealers have the same degree of risk aversion, and trades occur via publicly posted prices, then we get the more specific prediction that the two dealers with the most extreme inventories will be the first to trade. Given the difficulties mentioned above in defining the desired inventory level for dealers in unclassified cycles, we exclude these cycles from the analyses in this subsection.

Because we do not know that dealers have the same degree of risk aversion, we first report a very weak test of the inventory risk-sharing hypothesis. In Table VI we report the fraction of all IMM and IDB trades in which the buyer and the seller in the interdealer trade are trading to reduce (ex ante in the anticipated trades and ex post in unanticipated trades) their inventory risk. As a benchmark, if interdealer trade was independent of dealers' inventory positions, we would by *chance* expect to see 25 percent of all

²⁴ When we calculate these fractions for ambiguous cycles, we find that slightly less than 50 percent of IMM trades and nearly 60 percent of IDB trades unwind a dealer's extreme position. Thus, for ambiguous cycles there is no clear pattern. This is not too surprising, however, if these cycles include misclassified cycles or are ones where we cannot accurately distinguish between anticipated and unanticipated trades. Further, to the extent that ambiguous cycles include speculative positions, zero inventory may no longer be the dealer's desired inventory target.

Table VI
Interdealer Trades in Which Both Dealers Trade
to Unwind Positions

This table reports results for all interdealer trades in which both dealers are in a customer trade-related inventory cycle. IDB refers to trades on one of the four interdealer broker systems. IMM refers to direct trades between dealers registered to quote prices for a security.

	FTSE Stocks		Non-FTSE stocks	
	Number of Trades	Percent of Trades Unwinding	Number of Trades	Percent of Trades Unwinding
IMM trades	3,919	61.9	1,880	66.1
IDB trades	4,921	66.5	1,882	74.4

interdealer trades reducing both parties' inventories. The table provides striking evidence that interdealer trades pair dealers who are swapping positions. Between 62 and 74 percent of interdealer trades during cycles caused by customer trades are in the correct direction. That the rate is slightly higher for IDB trades compared to IMM trades is not surprising since IDB trades are voluntary, whereas IMM trades can occur involuntarily when another dealer places an order with a dealer posting at the best bid or ask.

As a next step, we look at the relative positions of the trade counterparties. As a practical matter, there are many different ways we could describe the counterparties' positions. The most obvious way is to report dealers' median inventories just prior to the interdealer trade. We choose medians because they remove the effect of any outliers caused by data or cycle coding errors. Because the level of inventories alone does not tell us where the inventories stand in relation to other dealers' inventories, we also rank the dealers' inventories just prior to the trade. The dealer with the largest inventory is assigned the value of 100. The dealer with the smallest (usually negative) inventory is assigned an index of zero.

Table VII reports the average inventory and indices of buyers and sellers in interdealer trades by cycle. Again we condition on whether or not the dealers anticipate their position and we require both dealers to be in an inventory cycle. The results for FTSE stocks show that the absolute median inventories of buying and selling dealers range from 4.1 to 5.3 NMS and are exactly as a risk-sharing model would predict. Consider, for example, an IMM trade where the median buyer is long 5.3 NMS and the median seller also is long 5.2 NMS. Here both buyer and seller are long because the buyer anticipates one or more coming customer sell orders and the seller is trying to unwind a long position he has accumulated. Looking across the rows and columns we can see that the median inventories match the predictions in Hypothesis 3.

To gain some perspective on how much of an inventory position these trades reverse, recall that the median peak cycle inventory is in the range of four to seven NMS (see Table II). In our sample, the median FTSE IDB (IMM) trade is 0.7 (1.0) NMS. This means that if dealers were to unwind their entire

Table VII
Buyer and Seller Inventories Prior to Interdealer Trades

This table reports the median inventories and median inventory ranks of the buying and selling dealers for all interdealer trades during customer trade-related inventory cycles. A dealer's inventory index equals $100 \times (R - 1)/(N - 1)$, where R is the dealer's rank in the inventory distribution and $R = 1$ for the market maker with the lowest inventory and N is the number of dealers in the stock. We define *anticipated* inventory cycles as cycles where interdealer trading precedes customer trades and *unanticipated* cycles as cycles where interdealer trading follows customer trades. IDB refers to trades on one of the four interdealer broker systems. IMM refers to direct trades between dealers registered to quote prices for a security. NMS is normal market size.

Buyer Cycle and Seller Cycle	FTSE Stocks				Non-FTSE Stocks			
	Inventory (NMS units)		Index (%)		Inventory (NMS units)		Index (%)	
	Buyer	Seller	Buyer	Seller	Buyer	Seller	Buyer	Seller
Buyer unanticipated short and seller unanticipated long								
IMM trades	-5.1	4.7	8	92	-5.5	5.6	8	100
IDB trades	-4.8	5.1	8	93	-5.5	5.7	0	100
Buyer anticipated long and seller anticipated short								
IMM trades	4.7	-4.8	86	15	4.3	-5.9	82	12
IDB trades	5.1	-5.3	86	13	5.8	-6.1	78	24
Buyer anticipated long and seller unanticipated long								
IMM trades	5.3	5.2	79	92	5.5	5.6	75	88
IDB trades	5.0	5.3	79	92	6.2	7.3	71	90
Buyer unanticipated short and seller anticipated short								
IMM trades	-5.0	-4.8	13	23	-5.7	-5.9	13	29
IDB trades	-5.3	-4.8	8	23	-6.1	-6.3	14	25

position with interdealer trades, they might require four to seven trades. According to Table III, because they get less than half their volume from customer trades, dealers will typically need between two and three interdealer trades to return to zero inventory.

Table VII also shows that the ranks of both buyers and sellers are very close to the extremes of the inventory distribution at the time of trade. For example, consider an IMM trade between a buyer who is positioning an unanticipated customer buy order and a seller who is positioning an unanticipated customer sell order in a FTSE stock with fourteen dealers. The median buyer has the second lowest inventory ($100 \times (2 - 1)/13 = 8$) and the median seller the second highest inventory ($100 \times (13 - 1)/13 = 92$) of all dealers in the stock. When one or both dealers are working an anticipated customer order, their inventories are generally somewhat less extreme. In the 14-dealer example, the median buyer and seller are the second or the third dealer from each end of the inventory distribution. Results for the non-FTSE sample are very similar.

The results in this subsection show that when dealers experience inventory imbalances induced by customer trades, they usually shift this imbalance to dealers who have opposite positions. Moreover, the dealers who trade with each other tend to be those at the extremes of the inventory distribution. Taken together, these results reject the null that there is no relationship between interdealer trades and the inventory distribution in favor of Hypothesis 3: London dealers use interdealer trades to share inventory risk from large customer trades.

D. Is Position Taking Profitable?

According to our data, more than two-thirds of our cycles contained at least one large trade exceeding three NMS. Because London dealers are not obliged to accept a three-NMS order, the fact that they do suggests that they must be compensated for the inventory risk they assume. Dealer revenue in principle consists of both trading revenue and commissions. We unfortunately do not have access to trade-by-trade commissions and therefore have to resort to estimating only trading revenue. For 40 percent of customer volume traded by dealers, this is the only source of revenue according to a recent study by Hansch et al. (1997). Based on Exchange estimates, we conjecture that another 15 percent (0.25×60) is likely to be commission-free (see Section I). Thus, dealers receive commissions from less than half of their customer order flow. Moreover, large trades are more likely to receive commission discounts or be commission-free. Thus, commission revenue is likely to be modest if charged at all on large customer trades. Since commission-free interdealer trades play such an important role in inventory management, we expect any risk premium for taking large positions to be reflected in dealers' trading revenues.

To test this hypothesis, we calculate trading revenues for each dealer and each cycle in the sample. Since cycles start and end at the same inventory level, we simply sum the pounds received by the dealer in stock sales during the cycle and deduct the pounds paid for the same number of shares. Unlike Hansch

Table VIII
Median Revenue for Dealers in Inventory Cycles

We calculate revenues by netting pounds received (in stock sales) against pounds paid (in stock purchases) by a dealer during a cycle. We calculate basis point revenues by dividing the pound revenue of each cycle by the value of shares bought and sold during the cycle and multiplying the result by 10,000. The number of cycles studied here is smaller than in previous tables because we eliminate cycles that have significant over-the-counter options trades. We do this because over-the-counter options prices differ from the prices for ordinary shares. We define *anticipated* inventory cycles as cycles where interdealer trading precedes customer trades and *unanticipated* cycles as ones where interdealer trading follows customer trades.

Cycle Type	Number of Cycles	Estimated Median Revenues	
		in Pounds	in Basis Points
FTSE stocks			
Unanticipated customer	5,430	1,098.61	11.6
Anticipated customer	2,898	1,418.92	15.4
Unanticipated broker	80	1,061.91	12.5
Anticipated broker	40	751.43	7.6
Crosses	382	1,264.13	17.6
Non-FTSE stocks			
Unanticipated customer	3,398	859.45	18.4
Anticipated customer	1,737	940.70	19.6
Unanticipated broker	57	727.55	18.9
Anticipated broker	49	931.51	21.2
Crosses	573	760.06	19.3

and Neuberger (1995) and Hansch et al. (1997), we include all trades executed by a dealer for his own account during the cycle to assess trading revenues. By contrast, Hansch et al. (1997) study the profitability of a *subset* of trades between dealers and customers, and explicitly exclude interdealer trades.

Table VIII reports the median revenue by cycle expressed in pounds and in basis points. We define the basis point revenue as 10,000 times the pound revenue for a cycle divided by the sum of the value of purchases and the value of sales executed by the dealer during the cycle. The results show that trading revenues are small but nonetheless significantly positive for all types of cycles resulting from block trades. Our median estimates range from £751 to £1,418 (8 to 15 basis points) per cycle for FTSE stocks and from £728 to £941 (18 to 21 basis points) per cycle for non-FTSE stocks. Median revenues are systematically higher for anticipated customer trades for FTSE stocks. By contrast, revenues are not systematically higher for anticipated customer trades for non-FTSE stocks. It is likely that this is a result of the lower liquidity of IDB systems for non-FTSE stocks (see Table I). There is relatively more heterogeneity in the revenues for broker trade-related cycles than for customer trade-related cycles, but we attribute this mostly to the small sample size. Taken together, these findings suggest that London dealers are rewarded for the risk associated with taking large customer (or broker) trades as we predict in Hypothesis 4.

We also report the median revenues for cycles where dealers cross offsetting large customer trades. Pound revenues from crosses are comparable to those caused by large customer trades. Since cycles resulting from crosses are extremely short-lived (see Table II), it seems that dealers are handsomely rewarded for matching large customer trades.²⁵

The results reported in Table VIII show that London dealers typically earn net positive trading revenues from all types of cycles. Of the dealers active in our sample, three (of seventeen) lose money in FTSE stocks and five (of sixteen) lose money in non-FTSE stocks. Thus, unlike Hansch and Neuberger (1995) who find negative revenues from position taking, we generally find positive trading revenues associated with position taking. We attribute the difference to the fact that Hansch and Neuberger synthetically create trading strategies from each dealer's order flow rather than studying the trading strategies actually used by London dealers.²⁶

Our final set of hypotheses relate to how interdealer trading affects trading revenues during customer trade-related cycles. If dealers are risk averse, then the larger their position, the greater the trading revenue required to compensate them for this risk.²⁷ Based on our earlier work, we also anticipate that there may be differences in the profitability of dealer positions depending on what types of interdealer trades the dealers use to unwind their positions. Specifically, earlier we find that IMM trades were more costly to execute than IDB trades.

To reduce the influence of outliers on our estimates, we estimate quantile regression models that relate the median revenues dealers earn over a cycle to various characteristics of the cycle. We only report regression results for FTSE stocks since the results for non-FTSE stocks are similar. The columns in Table IX report how cycle trading revenues vary for unanticipated and anticipated cycles. Panel A reports results for regressions that include controls for whether the cycle was completed during the same trading day, the absolute value of the peak inventory (in NMS units) minus three NMS, and the square of peak inventory (in NMS units) minus three NMS. The latter two variables ask whether broker-dealers are compensated for the extra risk associated with large positions. Asymptotic *t*-statistics are in parentheses below each estimated coefficient. The results show that trading revenues increase significantly with the dealer's maximum (in absolute value) position, but at a decreasing rate. A marginal increase in the peak of the cycle from 10 to 11 NMS would increase median profits by £183 ($191.0 - 2 \times$

²⁵ We find that ambiguous cycles are less profitable, both in pound terms and in terms of basis points, than customer trade-related cycles. For example, median trading revenues during ambiguous cycles for FTSE stocks are £308.50 or 2.3 basis points. One explanation for the lower profitability of the ambiguous cycles is that they last longer. Thus, dealers are exposed to greater price risk and ex post find these positions less profitable.

²⁶ Hansch et al. (1997) also decompose order flow to study the profits from certain types of trades.

²⁷ Although we test the null hypothesis that the size of the imbalance does not increase trading profits, one might expect our test to have low power since risk-averse dealers could extract a risk premium through their commissions.

Table IX

Quantile Regression Analysis of Cycle Revenues for FTSE Stocks

This table reports results from a quantile regression that explains variation in dealer trading revenue received during a customer trade-related inventory cycle. Same day dummy is a zero-one dummy equal to one if the inventory cycle was completed during the same day. Peak vol is the maximum position in normal market size (NMS) units that the dealer had during the cycle. Cust vol is the customer volume in NMS units in the direction opposite to that of the dealer's cycle position. We define interdealer broker (IDB), intermarketmaker (IMM), and broker vol similarly. The table reports estimated coefficients and asymptotic *t*-statistics. We pool data for long and short positions to conserve space. Each regression includes security fixed effects. We exclude cycles related to broker trades and ambiguous cycles. *Anticipated* inventory cycles are cycles where interdealer trading precedes customer trades and *unanticipated* cycles are ones where interdealer trading follows customer trades.

Independent Variable	Panel A		Panel B	
	Unanticipated Cycle Profit (£)	Anticipated Cycle Profit (£)	Unanticipated Cycle Profit (£)	Anticipated Cycle Profit (£)
Same day dummy	389.4 (3.8)	458.1 (4.6)	459.7 (5.3)	453.4 (4.8)
Abs(Peak vol - 3)	191.0 (23.2)	455.7 (58.42)		
(Peak vol - 3) ²	-0.55 (-20.6)	-0.57 (-57.2)		
Abs(Cust vol)			129.9 (19.8)	205.5 (31.9)
(Cust vol) ²			-0.38 (-16.0)	-0.23 (-30.3)
Abs(IMM vol)			-151.7 (-8.7)	-27.0 (-1.3)
(IMM vol) ²			7.4 (15.3)	32.0 (46.6)
Abs(IDB vol)			193.6 (11.7)	167.1 (7.5)
(IDB vol) ²			6.9 (15.5)	27.4 (28.7)
Abs(Broker vol)			250.3 (12.3)	-49.1 (-1.9)
Security fixed-effects included	Yes	Yes	Yes	Yes
<i>N</i>	5309	2807	5309	2807
Median Profit	540.5	2469.5	540.5	2469.5

$(10 - 3) \times 0.55$) for unanticipated cycles and £448 $(455.7 - 2 \times (10 - 3) \times 0.57)$ for anticipated cycles. Thus, it appears London dealers do get additional revenue from taking positions related to customer trades. We also find that there is an advantage to unwinding a position resulting from customer trades during the same trading day.

In Panel B we expand the analysis by including additional regressors corresponding to the value of each cycle that was managed using each of the four categories of trades: customer trades, IDB trades, IMM trades, and bro-

ker trades. In all but the case of broker trades, we include a squared term to capture nonlinearities. (We exclude the squared term for broker volume because there are few cycles with any broker volume.) Our relative cost argument suggests that cycles should be more profitable the more customer trades are used to manage the position. The estimates confirm this is the case. Moreover, we see that there are significant differences in trading revenues according to which type of interdealer trades are used to unwind positions. If the dealer is unwinding a position, it is expensive both in relative and absolute terms to use IMM trades. For example, a marginal IMM trade of one NMS costs the dealer a median £144 in trading revenues. Conversely, IDB trades are estimated to be much more desirable than either IMM or customer trades. A marginal one NMS IDB trade increases the dealer's revenues by about £200. Finally, broker trades appear to be the most profitable, although it appears that dealers have relatively few opportunities to trade with other brokers (Table III).

The results change somewhat when we examine customer trade-related anticipated inventory cycles. Now the dealers only earn higher trading revenue on marginal IDB trades than on customer trades if they are able to trade more than one NMS. Small IMM trades still result in lost trading revenues, large IMM trades generate limited trading revenues. Broker trades now are marginally unprofitable.

The result that IDB trades, and under some circumstances IMM trades, increase trading revenues on the margin might seem puzzling given our previous discussion. The profitability of IDB trades relative to IMM trades suggests that there are *systematic differences* in the conditions under which IDB and IMM trades are used for inventory management. Since IMM trades are more costly, we expect dealers to use these trades primarily as a last resort. Examples of such situations include when there is no liquidity in the four IDB systems and when the dealer must rapidly build or unwind a position. That IMM trades are less profitable than IDB trades is therefore not surprising. More puzzling is the fact that IDB trades consistently increase revenues. One reason might be that dealers primarily use IDB systems when they handle uninformative order flow. Alternatively, it is possible that dealers find that the anonymity of IDB systems reduces the price impact of trades on their own account. We find these differences between IDB and IMM trades intriguing and are exploring them more fully in related work (Reiss and Werner (1997a, 1997b)).

V. Conclusion

We have shown that interdealer trade is prevalent and facilitates risk sharing among London dealers. Our analysis first documents the extent of interdealer trading and its frequency during inventory imbalances. We then show that the pattern of interdealer trade is consistent with theoretical predictions from models of inventory risk sharing. For example, we report that more than 80 percent of interdealer trades during inventory cycles caused by

large customer trades are in the direction predicted by risk sharing. Our results also show that interdealer trades typically take place between the dealers with the most extreme inventory imbalances. Finally, we provide evidence that dealers are compensated for the risk associated with the large inventory positions they assume.

Our analysis also highlights how Exchange regulations and market making practices affect interdealer trade patterns. London dealers must be prepared to accept large block trades and it is difficult for them to attract (profitable) order flow by moving quotes because of preferencing arrangements. As a result, dealers frequently ask customers to give them time to position large orders (what we call *anticipated trades*) and they also make use of third-party interdealer brokers to trade anonymously with one another. Failure to account for these and other features of the London trading environment would make testing risk-sharing models of interdealer trade less meaningful.

The methodologies that we have developed for studying dealer inventories are but a first step toward understanding how dealer markets operate. We hope that future theoretical and empirical research on dealer markets will recognize the importance of interdealer trade for the functioning of dealer markets. In future work, we plan to pursue many open issues raised by our analysis. We plan to refine our method for classifying inventory cycles in ways that recognize differences among dealers. We also plan to explore the pricing, timing, and sequencing of interdealer trades. Finally, we would like to examine the extent to which dealers offset their inventory positions using other securities.

Appendix A

Our sample contains data for 1,923 of the most heavily traded U.K. and Irish equities during 1991. Each stock in the sample has an NMS greater than 2,000 shares. We do not have complete information because of problems the Exchange encountered in retrieving data. Specifically, the Exchange could not provide us with trade data for January 17, 18, and 21. The analysis includes all available data. Our results are robust to starting the sample on January 22, 1991.

Table AI compares the coverage of our data to figures published by the Quality of Markets Group at the London Stock Exchange. Our figures differ from the Exchange's figures for at least three reasons. First, as described above, we do not have all transactions. Second, we have removed canceled and corrected trades. Prior to the mid-1990s, the Exchange included canceled trades in its statistics. By our calculations, canceled trades account for approximately 2.4 percent (0.1 percent) of customer (interdealer) volume in U.K. and Irish equities and 2.5 percent (0.1 percent) by value of customer (interdealer) volume in the FTSE equities. Third, during 1991 the composition of the FTSE changed. When the Exchange reports annual statistics for the FTSE, they sometimes compute them assuming the year-end composi-

Table AI
1991 Quality of Markets Group Estimated London Trade Volume
by Type of Trade

The Exchange figures come from 1991 issues of the *Stock Exchange Quarterly*. The adjusted numbers are based on our sample of trades. The Exchange counts interdealer trades twice in its official statistics. The coverage fraction equals the ratio of our tabulated volume to Exchange estimates. N.A. means not available or not applicable.

Trade Type	U.K. and Irish Equities			FTSE Equities	
	Volume (million £)	Volume (%)	Coverage Fraction	Volume (million £)	Coverage Fraction
Reported by Exchange					
Intermarket	126,344	35.1	N.A.	N.A.	N.A.
Customer	234,116	64.9	N.A.	134,014	N.A.
Adjusted sample					
Intermarket	63,104	21.6	1.008	N.A.	N.A.
Customer	228,497	78.4	0.905	130,664	1.008

tion held for the entire year and they sometimes adjust for the changing composition. Our FTSE figures and the Exchange figures in Table AI adjust for the changing composition.

The adjusted trade figures at the bottom of Table AI correct the Exchange's double counting of interdealer trades. The Exchange counts interdealer broker trades as two trades: a sale from a dealer to the broker and a buy from the broker by another dealer. Because interdealer brokers do not trade for their own account, we treat interdealer trades as a single transaction between two dealers.


Appendix B

Inventories Created by Cumulating Trades

This appendix illustrates why we choose to exclude securities with significant overseas trading or securities that had close substitutes.

Many large British stocks also trade overseas. When U.K. dealers trade the underlying equity in London and depository receipts abroad, the dealer's position in London need not be related to the firm's worldwide position. One nice example we have of an overseas offset is Hansen PLC (HNSN). HNSN is a British stock and it also trades as a dollar denominated American Depository Receipt (ADR) in London. If one only used inventory data on the British stock HNSN, one would be led astray. For example, Figure B1 graphs the apparent inventories of a dealer in HNSN, who seemingly is accumulating a 15,000,000 share position during 1991. If we compute his position in the ADR, however, we find that the dealer is selling the ADR at the same

Panel A



Panel B




Figure B1. Hansen plc.—an ADR example. The figure displays the actual 1991 intraday inventory series for one dealer registered as a market maker both in Hansen’s underlying stock and in its London-traded ADR. Each trade for the dealer’s own account is cumulated starting from zero to create an inventory series represented by a solid line. The ADR inventory series is adjusted to underlying share units.

time. On balance, he has almost no inventory! Though an obvious solution to this problem would be to include dealers' trading in overseas instruments, the London data do not include all overseas trades of U.K. dealers.

The London market also has nearly identical equity issues trading at the same time. The most common example of such equities in London are traded rights or warrants. For example, on 5 July 1991 Sainsbury PLC (SBRY) issued rights to its stock. These rights traded separately for a time and then merged with the underlying stock when the rights issues were fully paid. Figure B2 illustrates how rights issues can create problems for inventory data. In Figure B2, a dealer in the underlying stock appears to have his inventory falling drastically after the rights issue. However, if we look at the dealer's position in the rights issue, we see that his position is nearly offset.

Our third example is Witan Investment Co. (WTAN). WTAN has an actively traded warrant during 1991. Figure B3 shows the inventory of a dealer who seemingly takes a 400,000 share short position in WTAN. However, the dealer is simultaneously taking a long position in the warrant. After valuing the warrants in terms of the underlying, we find that the dealer's warrant position comes close to offsetting his position in the underlying stock.


These examples and others lead us to choose stocks that (i) do not have liquid ADRs trading in London; (ii) do not have rights or seasoned equity issues; (iii) do not have actively traded derivative securities; and (iv) are not investment trusts. We exclude investment trusts because their inventory risk can potentially be hedged by holding the underlying securities. Table BI lists the stocks in our sample.

Inventory Cycles

Our inventory cycle algorithm cumulates each dealer's trades by stock. Specifically, the algorithm cumulates the number of shares bought and sold from different sources: customers and brokers, IDBs, and other dealers (IMM). By definition, the total number of shares purchased will equal the total number of shares sold. The algorithm then evaluates each inventory series, keeping track of when a dealer's inventory crosses the boundaries defined in Figure 2. If a cycle starts by the inventory crossing +1 NMS (Point 2) from below, and the inventory subsequently crosses +3 NMS (Point 3) from below without falling below +1 NMS, then the cycle is classified as a long position. We continue cumulating purchases and sales until the inventory crosses +1 NMS from above (Point 4). The algorithm for short cycles is symmetric.

We classify cycles as anticipated and unanticipated as follows. Once we have identified a long or short cycle, we calculate the total number of customer buys, interdealer buys, customer sales, and interdealer sales. We classify a long cycle as unanticipated if at least two-thirds of the sell volume is from customers selling to the dealer. (A similar definition applies to short cycles.) If less than two-thirds of the sell volume is from customers, then we check buy volumes. We classify the cycle as anticipated if it is not an un-

Panel A



Panel B




Figure B2. Sainsbury plc.—a rights example. The figure displays the actual 1991 intraday inventory series for one dealer registered as a market maker both in Sainsbury's underlying stock and in one of its rights issues. Each trade for the dealer's own account is cumulated starting from zero to create an inventory series represented by a solid line. The rights issue inventory series is adjusted to underlying share units.



Panel A**Panel B**

Figure B3. Witan Investment—a warrants example. The figure displays the actual 1991 intraday inventory series for one dealer registered as a market maker both in Witan's underlying stock and in warrants on Witan stock. Each trade for the dealer's own account is cumulated starting from zero to create an inventory series represented by a solid line. The warrants inventory series is adjusted to underlying share units.

Table BI
1991 Characteristics of Sample Stocks

This table lists the stocks in our sample. We include stocks that: (i) do not have liquid ADRs trading in London; (ii) do not have any rights or seasoned equity issues; (iii) do not have trades in other securities included in our data set; and (iv) are not investment trusts. These criteria leave us with twenty-five FTSE stocks. We sort eligible non-FTSE stocks by number of dealers. We then select the two stocks with the highest market capitalization for a given number of dealers. Market capitalizations, average trade prices, and normal market sizes (NMS) are as of 31 March 1991 and are from the *Quality of Markets Companies Book 1991*, Table 1. The EPIC code is the Exchange's mnemonic for security ticker symbols.

	EPIC Ticker	Market Capitalization (million £)	Average Price (£)	NMS (1,000 shares)	Maximum No. of Dealers
FTSE stocks					
Abey National	ANL	3,564.0	2.72	50	14
BAA	BAA	2,042.4	4.07	50	15
BICC	BICC	1,230.1	4.46	25	13
Blue Circle Industries	BCI	1,497.6	2.70	50	16
Commercial Union	CUAC	2,209.6	5.09	25	13
Enterprise Oil	ETP	2,681.6	5.88	25	16
Forte	FTE	2,121.4	2.67	50	15
General Electric Co.	GEC	5,424.2	2.01	100	15
Guardian Royal Exchange	GARD	1,863.2	2.16	50	13
Kingfisher	KGF	2,052.2	4.52	25	14
Land Securities	LAND	2,708.5	5.37	25	11
Legal & General Group	LGEM	2,242.2	4.65	25	13
Lloyds Bank	LLOY	4,171.2	3.36	50	15
MEPC	MEPC	1,677.1	5.20	15	9
Midland Bank	MID	1,567.3	2.00	50	15
Pilkington	PILK	1,409.5	1.88	50	13
RMC Group	RMC	1,347.7	6.94	10	14
Royal Insurance Holdings	ROYL	2,156.2	4.45	25	13
Scottish & Newcastle	SCTN	1,485.2	3.74	25	15
Sears	SEAR	1,324.5	0.88	75	14
Smith & Nephew	SN.	1,286.7	1.29	50	16
Sun Alliance Group	SUN	3,002.0	3.78	25	13
Tarmac	TARM	1,825.9	2.51	50	15
TSB Group	TSB	2,240.1	1.49	75	14
Arjo Wiggins Appleton	AWA	1,793.4	2.15	50	11

anticipated cycle and more than two-thirds of the buy volume is from customers. We have relatively few cycles where neither the customer buy or sell volume exceeds two-thirds of the total buy or sell volume.

Once we classify the cycle as either anticipated or unanticipated, we calculate the proportion of the total cycle managed using IDB, IMM, and customer trades. To illustrate, for a long unanticipated cycle, we calculated the fraction of IDB trading as

$$\frac{\text{MM Sales via IDBs}}{\text{MM Interdealer Buys} + \text{MM Buys from Customers}} \quad (\text{B1})$$

REFERENCES

- Biais, Bruno, 1993, Price formation and equilibrium liquidity in fragmented and centralized markets, *The Journal of Finance* 48, 157–184.
- Board, John, Anne Fremault Vila, and Charles Sutcliffe, 1997, Market maker heterogeneity: Evidence from the London Stock Exchange, Working paper, LSE Financial Markets Group.
- Board, John, and Charles Sutcliffe, 1995, The effects of trade transparency in the London Stock Exchange: A summary, London School of Economics Financial Markets Group Special Paper no. 67.
- Easley, David, and Maureen O'Hara, 1987, Price, trade size, and information in securities markets, *The Journal of Financial Economics* 19, 69–90.
- Foster, Douglas F., and S. Viswanathan, 1993, Variations in trading volume, volatility, and trading costs: Evidence on recent price formation models, *The Journal of Finance* 48, 187–211.
- Gould, John F., and Allan W. Kleidon, 1994, Market maker activity on NASDAQ: Implications for trading volume, *Stanford Journal of Law, Business & Finance* 1, 1–17.
- Greene, William H., 1997, *Econometric Analysis* (Prentice-Hall, Englewood Cliffs, N.J.).
- Hansch, Oliver, Narayan Naik, and S. Viswanathan, 1997, Preferencing, internalization, best execution and dealer profits, Working paper, London Business School.
- Hansch, Oliver, Narayan Naik, and S. Viswanathan, 1998, Do inventories matter in dealership markets? Some evidence from the London Stock Exchange, *The Journal of Finance* 53, 1623–1656.
- Hansch, Oliver, and Anthony Neuberger, 1995, Market maker profits on the London Stock Exchange, Working paper, London Business School.
- Hasbrouck, Joel, and George Sofianos, 1993, The trades of market makers, *The Journal of Finance* 48, 1565–1594.
- Ho, Thomas, and Hans R. Stoll, 1983, The dynamics of dealer markets under competition, *The Journal of Finance* 38, 1053–1074.
- Kyle, Albert S., 1989, Informed speculation with imperfect competition, *The Review of Economic Studies* 56, 317–356.
- Lyons, Richard K., 1995, Tests of microstructural hypotheses in the foreign exchange markets, *The Journal of Financial Economics* 39, 321–351.
- Lyons, Richard K., 1996, Optimal transparency in a dealership market with an application to foreign exchange, *The Journal of Financial Intermediation* 5, 225–254.
- Madhavan, Ananth, and Seymour Smidt, 1993, An analysis of daily changes in specialist inventories and quotations, *The Journal of Finance* 48, 1595–1628.
- Manaster, Steven, and Steven C. Mann, 1996, Life in the pits: Competitive market making and inventory control, *The Review of Financial Studies* 9, 953–975.
- Naik, Narayan, Anthony Neuberger, and S. Viswanathan, 1996, Disclosure regulation in competitive dealership markets, Working paper, London Business School.
- Neuberger, Anthony, and Robert A. Schwartz, 1989, Current developments in the London equity market, Working paper, New York University.

- New Electronic Trading Service*, 1996, London Stock Exchange consultation document (London Stock Exchange, London, U.K.).
- O'Hara, Maureen, 1995, *Market Microstructure Theory* (Basil Blackwell Business, Cambridge, MA).
- Oliver, Wyman & Co, 1997, *Review of block trading and publication requirements for UK stocks*, London Stock Exchange commissioned report, April 1997.
- Reiss, Peter C., and Ingrid M. Werner, 1995, Transaction costs in dealer markets: Evidence from the London Stock Exchange, in Andrew Lo, ed.: *The Industrial Organization and Regulation of the Securities Industry* (University of Chicago Press, Chicago, Ill.).
- Reiss, Peter C., and Ingrid M. Werner, 1997a, Friend or foe? The pricing of London interdealer trades, Mimeo, Stanford Business School.
- Reiss, Peter C., and Ingrid M. Werner, 1997b, To split or not to split? Managing the price impact of block trades, Mimeo, Stanford Business School.
- Röell, Ailsa, 1992, Stock market transparency, Working paper, London School of Economics.
- Rules for the Stock Exchange Electronic Trading Service*, 1997, Stock Exchange Notice No. 4 (London Stock Exchange, London, U.K.).
- Stock Exchange Quarterly*, 1991, Winter Edition (London Stock Exchange, London, U.K.).
- Stoll, Hans R., 1976, Dealer inventory behavior: An empirical investigation of NASDAQ stocks, *The Journal of Financial and Quantitative Analysis* 11, 359–380.
- Tonks, Ian, and David Webb, 1989, The reorganization of the London stock market: The causes and consequences of “Big-Bang,” London School of Economics Financial Markets Group Special Paper no. 20.
- Vayanos, Dimitri, 1998, Strategic trading and welfare in a dynamic market, *The Review of Financial Studies* 11, 1–58.
- Vogler, Karl-Hubert, 1993, Interdealer trading, London School of Economics Financial Markets Group Discussion Paper no. 174.
- Werner, Ingrid M., 1997, A double auction model of interdealer trading, Stanford Business School Research Paper no. 1454.