Bund for Glory, or It's a Long Way to Tip a Market

Craig Pirrong
University of Houston
spirrong@uh.edu
713-743-4466

March 8, 2005

Abstract. Theory predicts that liquidity considerations make financial markets "tippy." In 1998, trading on Bund futures tipped from LIFFE (an open outcry exchange) to Eurex (an electronic market). Measures of spreads on LIFFE and Eurex did not change markedly in the eighteen month period over which Eurex achieved dominance in Bund futures trading, but a measure of market depth did worsen on LIFFE as tipping proceeded. The evidence suggests that trading fee differentials and operational efficiences were the key factors in preciptiating the shift in volume. The "sponsorship" of the Eurex platform by German banks narrowed liquidity cost differences sufficiently to permit Eurex to charge lower fees and thereby undercut total trading costs on LIFFE.

JEL Classification: L11, L12, L31, G10, G20. Key Words: Securities market structure, financial exchanges.

1 Introduction

Theory predicts that financial markets are "tippy," that is, that all trading volume tends to gravitate to a single trading venue (Admati and Pfleiderer, 1988, 1991; Pagano, 1989; Grossman, 1992; Glosten, 1994; Madhavan, 2000; Pirrong, 2002, 2003a-c). In the presence of informed traders, liquidity traders minimize their transactions costs by congregating in a single market. Theory predicts that liquidity considerations create network effects that make trading in a particular instrument a natural monopoly.[1]

This article analyzes empirically a well-known episode of tipping–the movement of trading in Bund futures (one of the most heavily traded financial instruments in the world) from the open outcry London International Financial Futures and Options Exchange ("LIFFE") to the electronic Eurex. At the beginning of 1997, about 65 percent of Bund futures trading took place on LIFFE. Over the next 21 months all trading volume switched to Eurex. Eurex has maintained its 100 percent market share of Bund futures trading, and also has virtually 100 percent of the trading of futures on other German government debt instruments.

I document several interesting results.

- Throughout the January, 1997-June, 1998 period there were only small differences measures of bid-ask spreads on LIFFE and Eurex, with the

[1]Most cases of market fragmentation–the trading of a given instrument in multiple venues–are attributable to "cream skimming." That is, most satellite markets survivie only by restricting trading to the verifiably uninformed. See Pirrong (2002, 2003b) for a formal analysis predicting this result and a discussion of the extensive empirical evidence supporting this prediction.

Roll spread measure being somewhat smaller on LIFFE.

- Despite the fact that LIFFE steadily lost market share and Eurex steadily gained market share from January, 1997 to August, 1998, there was no pronounced rise in spread measures on LIFFE and no pronounced fall in spread measures on Eurex until LIFFE's share of volume had fallen to less than 18 percent in June, 1998.

- Measures of price impact (depth) provide evidence of more marked changes in liquidity on the two exchanges beginning in mid-1997, and accelerating in early 1998. Although the price impact of trades on LIFFE was considerably smaller than on Eurex in mid-1997, Eurex had achieved rough parity in price impact by November, 1997. Moreover, Eurex had a decisive advantage in depth as early as March, 1998, months before the deterioration of LIFFE spreads.

- Most important, the events of 1997-1998 suggest that exchange pricing decisions and non-liquidity-related costs can decisively influence which trading venue prevails in a competitive contest. Starting in 1997, Eurex and LIFFE implemented a series of competitive price cuts. LIFFE's loss of market share accelerated markedly when Eurex slashed fees in January, 1998, and LIFFE did not respond immediately. In the aftermath of the Eurex price cut, the sum of exchange fees and liquidity costs (as measured by price impact) was lower on the German exchange; the pace of "tipping" accelerated noticeably at this time. Furthermore, best available estimates suggest that the labor and equipment costs of trading on an automated exchange are approximately one-third as large

> as the costs of trading on an open outcry market. Thus, all in trading costs–including exchange fees and brokerage costs in addition to spread and depth related trading costs–are relevant in determining when and where trading tips. In this one particular instance, trading tipped to the market that could process trades more cheaply; the trade processing efficiencies on Eurex were sufficiently large to overcome LIFFE's initial liquidity advantage. This is consistent with the implications of Pirrong (2002, 2003b).

Eurex's operational efficiencies and fee cuts were decisive because liquidity cost differences between Eurex and LIFFE were sufficiently small (as measured empirically) by late-1997. The support of the German banks in the years 1990-1997 helped Eurex generate the trading volume necessary to narrow these liquidity cost differences. Thus, the LIFFE-Eurex Bund episode suggests that operational efficiencies are not sufficient to achieve tipping. Just building a more efficient platform does not ensure they will come; "sponsorship" by those in control of large order flows is important as well.

Although this article obviously has implications regarding the nature of competition between financial exchanges, it is also of interest for the study of network industries generally. Other "new economy" industries (e.g., computer software) are widely considered to be "tippy," but it is difficult to measure the costs and benefits of network effects, because many of these (e.g., the benefits of having more programs for a particular operating system) are difficult to quantify. In contrast, in the LIFFE-Eurex case it is possible to quantify with some precision the cost differences across "platforms" subject to network economies. The results demonstrate that very small cost differences (on the order of 3 percent of total costs) are sufficient to induce tipping in a market subject to network economies. The Eurex-LIFFE episode also highlights the importance of technology "sponsorship" in a network industry.

The remainder of this paper is organized as follows. Section 2 presents a brief history of the Bund futures market. Section 3 describes the empirical methodology. Section 4 presents the results. Section 5 summarizes the main implications of the analysis.

2 A History of Bund Futures Trading

Trading of futures on ten-year German government bonds,"Bunds," commenced in 1988, six years after the the formation of LIFFE. LIFFE was an open outcry futures exchange consciously modeled on the Chicago exchanges. Bund futures volume on LIFFE grew substantially from the contract's debut. It soon became the exchange's largest contract.

Eurex was created by the 1998 merger of the Swiss Exchange SOFFEX and the German Deutche Terminborse ("DTB").[2] DTB was formed in 1990 by a consortium of German banks. The exchange launched a Bund futures contract in competition with LIFFE. By 1993, DTB's Bund market share had grown to 24 percent. Most of this volume came from German banks, whereas most of the London and US originated order flow was directed to LIFFE.

[2]Originally, SOFFEX and DTB formed Eurex as a joint venture in 1996. The two exchanges announced a merger in September, 1997, and the merger was completed in 1998. Hereafter, I will call the German exchange "Eurex" regardless of the time to which I refer.

In 1996, Eurex launched an aggressive campaign to win market share from LIFFE. By December, 1997, Eurex's market share had grown to 40 percent. Key events in the rivalry between Eurex and LIFFE include:

- March, 1996. DTB provided screen-based access to its market from London.

- March, 1997. The United States futures regulator, the Commodity Futures Trading Commission, granted Eurex permission to install trading terminals in the United States. By October, 1997, ten futures commission merchants in the US had installed terminals, and accounted for 18 percent of Eurex Bund futures volume (Young and Theys, 1999).

- 1 August, 1997. Eurex extended its trading hours to match LIFFE's.

- 1 September, 1997. Eurex commenced trading fee waivers on the Bund. The fee waivers expired on 31 December, 1997. Beginning on 18 September, 1997, LIFFE also waived fees for the remainder of the year.

- 1 January, 1998. Eurex introduced a new pricing structure. It eliminated admission and annual fees for market makers and full members.[3] It also reduced fixed fees for clearing members. Furthermore, Eurex introduced a new fee structure that capped the per contract charge at .50 DM (the pre-September, 1997 price) and introduced a sliding

[3]These fees were substantial. A non-clearing membership required payment of a one-time DM 102,000 entry fee, a one-time line fee of DM 25,000, a monthly line fee of approximately DM 4000, and an annual fee of DM 34,000.

pricing scale that gave volume discounts. The resulting price schedule eliminated line fees and capped a user's total trading fees at 4500 DM (Young and Theys, 1999). This effectively set the marginal trading fee equal to zero for medium and large traders.

- 27 March, 1998. LIFFE reduced transaction fees for all of its financial futures and options. Fees for futures trades fell from 42 pence per contract to 25 pence per contract.

Eurex's efforts, and LIFFE's belated response, led to a pronounced shift in volume from LIFFE to Eurex. Figure 1 depicts the Bund futures market shares of the two exchanges over the January, 1997-October, 1998 period.[4] Over this period, LIFFE market share fell steadily. In January, 1998 Eurex volume surpassed LIFFE's. As Figure 1 makes clear, thereafter the erosion of LIFFE market share accelerated. From January to December, 1997, LIFFE market share fell about 1.1 percentage points per month; From January to July, 1998, LIFFE's market share fell an average of 8 percentage points per month. By July, 1998, Eurex's market share was in excess of 95 percent; by October of that year, the German exchange's market share exceeded 99.9 percent. The market had tipped completely.

The acceleration in market share changes once Eurex had surpassed a 50 percent share is broadly consistent with received theories on equilibrium in network industries, and with models of liquidity in financial markets. These theories imply that, *ceteris paribus,* trading costs are lower in the market

[4]These figures are based on data published in Eurex Monthly Statistics (1998), available at http://www.eurexchange.com and data on LIFFE Bund volume obtained from the Commodity Research Bureau.

with the greater market share. This lower cost attracts additional volume to that market, which exacerbates the cost disparity, which in turn leads to further movement of trading volume to the larger exchange. This process ends when one exchange captures all the volume.

Although the market share movements are illuminating, a more direct test of these theories requires a comparison of trading costs and market quality on the two exchanges. Moreover, a more detailed analysis of the data is required to understand why the market tipped when it did. The remainder of this article utilizes data from the two exchanges to (a) document the impact of market tipping on conventional measures of trading cost and market quality, and (b) determine what caused the market to tip in January, 1998.

3 Empirical Methodology

3.1 Data

The analysis is based on tick data obtained from LIFFE and Eurex. The LIFFE data includes the time and price of each trade, as well as bid and ask quotes (and their times). The Eurex data reports only time and price of trades; it does not report quotes. Although LIFFE data purports to report the volume of each trade, these figures are suspect. For instance, a large number of trades report a volume of 41 contracts. The Eurex data report the volume of each transaction, but do not report trade direction (i.e., whether a trade was buyer- or seller-initiated).

3.2 Spreads

Spreads are the most commonly employed measure of trading costs. I use several well-known spread measures to estimate trading costs on LIFFE and Eurex during the 1997-1998 period. Some measures require the observation of the quoted spread, and hence I can only implement them for LIFFE. Other measures do not require knowledge of the quoted spread and can be implemented for both exchanges.

The effective spread and the realized spread require knowledge of the quoted spread. The effective spread at time t is defined as:

$$S_t^E = 2Q_t(P_t - M_t)$$

where P_t is the transaction price at t, M_t is the midpoint of the quoted bid-ask spread observed immediately prior to the transaction, and Q_t is a trade direction indicator. The realized spread at t is:

$$S_t^R = 2Q_t(P_t - M_{t+.25})$$

where $M_{t+.25}$ is the mid-point of the quoted spread 15 minutes (a quarter hour) after t.

Transactions are included in the estimates of the effective and realized spreads only if there are both a bid price and an ask price time stamped within 15 seconds prior to the time of the transaction.

The Roll spread does not require knowledge of the quoted spread. The Roll spread is:

$$S^{Roll} = \sqrt{-2cov(\Delta P_t, \Delta P_{t-1})}$$

where $\Delta P_t = P_t - P_{t-1}$, and $t-1$ is the time of the transaction immediately preceeding the transaction at t.

Stoll's traded spread measure can also be implemented without knowledge of the quoted spread. This spread is the coefficient S^T estimated from the following regression:

$$\Delta P_t = .5S^T[Q_t + Q_{t-1}(\gamma - 1)] + \xi_t$$

In this expression, γ measures the fraction of the traded spread attributable to adverse selection and inventory costs. $1-\gamma$ measures market maker rents. ξ_t is an error term.

The Thompson-Waller measure also does not require knowledge of the quoted spread. This measure is:

$$S^{TW} = E[|P_t - P_{t-1}|]$$

where this expectation is calculated conditional on $P_t \neq P_{t-1}$.

All spread measures except the Roll and Thompson-Waller estimators require observation of the trade direction indicator Q_t. For LIFFE, where the bid-ask is observable, I use the Lee-Ready trade direction indicator.[5] For Eurex data, I estimate Stoll's traded spread measure using the tick test to determine Q_t; to permit comparison, I estimate the relevant regression on LIFFE data using the tick test-based trade direction measure, although I also use the Lee-Ready measure for LIFFE data. Given that Stoll shows that the possibility that broken up orders can distort the estimates of S^T

[5]The Ellis-Michaely-O'Hara (2000) trade indicator is almost always equal to the Lee-Ready indicator, and generates identical empirical results.

and γ, I estimate the regressions on a subset of data in which trades on a zero tick are eliminated as well as on the whole data. For brevity, only the results from the former analysis are presented; the conclusions do not depend on this choice.

I estimate each spread measure for each month beginning in January, 1997 and ending in August, 1998. Reported effective, realized, and Thompson-Waller spreads are sample averages of the relevant statistic within the month. The reported Roll spread is based on the within-month covariance of consecutive price changes (with overnight price changes omitted). The traded spread model is estimated month-by-month using GMM with Q_t and Q_{t-1} as instruments; overnight observations are omitted. All spread measures in a given month are based on data from the next-to-expire (i.e, "front month") contract. Contract months are rolled at the beginning of each delivery month. Thus, for example, in March (a delivery month) spreads are estimated using June futures prices, whereas in February the March prices are used.

3.3 Price Impact and Depth

Market depth measures how buy-sell volume imbalances impact prices. Extending Hasbrouck (2002), I implement a Markov Chain Monte Carlo method for estimating the price impact of transactions. Hasbrouck's method assumes that the volume of each transaction is available, but that trade direction (whether a particular trade is seller-initiated or buyer-initiated) is latent. Given the price and volume information, Hasbrouck's MCMC method draws inferences about trade direction and price impact through simulation and repeated application of Bayes' Law.

The Eurex data reports trade size, but not trade direction. As noted above, the LIFFE data do not contain reliable measures of trade size. Since trade-by-trade volume is available for Eurex, Hasbrouck's method can be implemented for this exchange. Such data is not available for LIFFE, however, so it is necessary to treat trade size (not just trade direction) as latent too. This requires a modification of Hasbrouck's method.

The Appendix contains a detailed description of the latent volume method, and Hasbrouck (2002) presents a detailed description of the latent trade direction method. Here I merely touch on the underlying motivation and implementation of these methods in broad strokes.

The underlying price determination model is based on fundamental microstructural considerations which imply that transactions can have permanent and transitory impacts on prices. Information effects lead to permanent price changes. Non-informational frictions, such as inventory or trade processing costs, lead to transitory effects ("price reversals").

An extension of the microstructure model of Brown-Zhang (1997–"BZ" hereafter) incorporates both temporary and persistent trade impacts.[6] In the BZ model, the price impact of a trade of size V is broken into two parts, one due to the informational content of the trade, the other due to an inventory effect.

Consider the following representation of an empirical microstructure model in the spirit of BZ. Index transactions by $t = 1, 2, \ldots, T$. Define V_t as the volume of transaction t (with positive values indicating a purchase and a neg-

[6]This model essentially extends the classical Kyle model by allowing risk averse market makers and hence inventory effects.

ative value indicating a sale), m_t as the "true" log value of the instrument at the time of transaction t, and τ_t the time elapsed since the last transaction prior to trade t. The true value of the instrument is its value conditional on publicly available information and current and past transactions.

The natural logarithm of the true value of the instrument follows:

$$m_t = m_{t-1} + \lambda V_t + \varepsilon_t \sqrt{\tau_t} \tag{1}$$

Here ε_t is a disturbance that reflects the flow of public information relating to true value that occurs between trades $t-1$ and t. The rate of public information flow is assumed constant over time, and hence ε_t has constant variance σ^2; scaling ε_t by $\sqrt{\tau_t}$ reflects the fact that transactions are not necessarily spaced evenly in time.[7] Note that trade volume impacts true value, hence λ estimates the information-related price impact per unit of volume. Also note that since m_{t-1} appears in the expression, λ reflects the permanent price impact of the trade. λ therefore corresponds to the first term in BZ expression (9), and represents the price impact of a trade attributable to private information.

Consider first price dynamics assuming that prices are continuous–that

[7]The scaling of the residual by the square root of the time since the last trade implies that the log price process is heteroskedastic, but that the form of the heteroskedasticity is known. Although some implementations of MCMC methods to microstructure problems do not take this scaling into account (and hence assume that trades are equally spaced in time), such implementations are mis-specified. The change in price between two trades that are very close in time is almost exclusively due to the price impact of the transactions, rather than the flow of public information over this time interval. In contrast, the flow of public information has a greater influence on price movements that occur between trades widely separated in time (e.g., between the close and the succeeding open). When quantifying trade impact (permanent and temporary) nearly contemporaneous trades should be given more weight than trades that occur some time apart.

is, assuming that there is no finite tick. In this case, the natural logarithm of the price of transaction t is:

$$p_t = m_t + \beta V_t. \tag{2}$$

The parameter β reflects the non-informational price impact per unit volume. Note that this impact is not persistent. It is analogous to the second term in BZ expression (9), and reflects the trade impact associated with inventory effects.[8]

Together, (1) and (2) imply:

$$\Delta p_t \equiv p_t - p_{t-1} = \lambda V_t + \beta(V_t - V_{t-1}) + \varepsilon_t \sqrt{\tau_t} \tag{3}$$

For Eurex, volume data–which corresponds to $|V_t|$–is available, but trade direction is not. Defining the direction of trade t as δ_t, with $\delta_t = 1$ for a buy and $\delta_t = -1$ for a sale, $V_t = \delta_t |V_t|$. The δ_t are latent, but MCMC methods can be used to infer the δ_t and the parameters of interest (λ, β, and σ). The details are contained in Hasbrouck, but the basic "recipe" is as follows:

1. Make a guess for the δ_t and select priors for the λ, β, and σ. I use almost uninformative priors. That is, the priors are uninformative except that they constrain λ and β to be non-negative, which means that I assume that buys (sells) cause prices to increase (decrease); this is reasonable *a priori,* and is necessary to identify the model.

2. Given the priors, the δ_t, and the volume and price data, use (3) to estimate the posterior densities for λ, β, and σ, and then draw from these densities.

[8] A more complicated model can incorporate trade processing costs as well.

3. The draws for λ, β, and σ, and the price and volume data imply posterior probability distributions for each of the δ_t–see Hasbrouck (2002) for their derivation. Draw the δ_t from these posteriors.

4. Return to step 2, and repeat S times. Discard the first S_1 "burn in" draws, and average the remaining $S - S_1$ draws for λ, β, and σ. These averages are estimates of the mean values for these parameters.

Accounting for discrete tick size greatly complicates the analysis. In this case, changes in the "true" prices are not observed, so (3) cannot be estimated. Instead, the econometrician only knows that if a trade is a buy at the price P_t, $m_t \in [\ln(P_t) - \beta V_t, \ln(P_t - \theta) - \beta V_t]$ where θ is the tick size.[9] In words, this means that price is rounded up to the nearest tick for a buy. Similarly, price is rounded down to the lowest tick for a sell.

This information can be used to make draws for m_t. Given the draws for m_t and V_t, (1) can be used to estimate the posterior density for λ, from which a draw can be made. With discreteness one cannot draw β simultaneously with λ. Instead, I utilize a Random Walk Metropolis Hastings method to estimate β.[10]

[9] A similar analysis holds for a sell–see the analysis in the appendix.

[10] As with Hasbrouck (2002), I find that tick size is sufficiently large compared to the the typical inventory effect that the estimated β is typically very close to zero. In the results reported below I impose $\beta = 0$. This greatly speeds computation, and experimentation indicates that this has little effect on the estimated λ's. I have also experimented with models that include a buy-sell dummy in (1), and with models that allow for a concave relation between trade impact and the size of a trade. The coefficient on the buy-sell dummy is typically very close to zero. Indeed, I constrain draws on this coefficient to be non-negative, and the density of the draws is concentrated around zero. The results reported below obtain for all specficiations examined.

For LIFFE data, not even $|V_t|$ is observed. Hence, a more complicated method that draws V_t, not just δ_t, is required. The appendix sets out this method in detail. The main implementation difficulty here is identification. Note that (2) is not identified absent additional information. Doubling all the volumes and halving λ and β would produce the same value for the likelihood function. To ensure that volumes drawn by the method reflect actual volumes in the marketplace, I estimate the model under the restriction that the absolute volumes drawn at each iteration of the MCMC sampler for a given trading day add up to the observed trading volume on that trading day. Such daily volume figures are available for LIFFE.

Extensive experimentation indicates that λ coefficients estimated without actual data on $|V_t|$ are somewhat larger than those estimated when such data is available, so direct comparisons of price impacts estimated using the two methods are problematic. However, these experiments indicate that the λ coefficients estimated using the two methods are highly and positively correlated.[11]

Therefore, I utilize several estimates of λ. First, I estimate this parameter using Eurex price and volume data assuming that only trade direction is latent. Second, I estimate this parameter using Eurex price data and assuming that volume is latent. Third, I estimate this parameter using LIFFE price data and assuming that volume is latent. The latter two results are

[11]The experiment involves estimating λ using both methods for markets where both price and volume data are available. For instance, I estimate λ using Eurex data assuming that only trade direction is latent, and then estimate it again for the same time periods assuming that volume is latent. Volume and price data is also available for sevaral Chicago Mercantile Exchange contracts, so I have executed the experiment using this data as well.

directly comparable, permitting inferences about relative depths at LIFFE and Eurex. All three estimates permit me to draw inferences about changes in market depth over time.

It should be noted that his method is very computationally intensive. By mid-1998, more than 10,000 trades occurred each day on Eurex. Each trade must be simulated S times (I use $S = 2000$ and $S_1 = 400$). Each simulation involves draws from probability distributions. Estimation of the relevant coefficient for a single day's trading can take upwards of three hours on a 1.2 GHZ machine.

To economize on computation cost, I estimate the λ coefficient for one day (Wednesday) of each week beginning in June, 1997 and ending in July, 1998. This produces 60 λ coefficients for each of the three estimations (Eurex with volume, Eurex without volume, and LIFFE without volume).

4 Results

4.1 Spreads

Tables 1 and 2 present the results for the spread analysis. Table 1 presents results on the traded spread, Thompson-Waller, and Roll spread measures, which can be compared across exchanges. Table 2 presents results on the spread measures that are available only for LIFFE due to the absence of bid-ask quotes on Eurex. In each table, the first column denotes the contract month.[12] The second column denotes the calendar month.

[12] H, M, U, and Z are standard futures industry delivery month codes for March, June, September, and December, respectively.

Several regularities stand out in Table 1:

- With respect to the traded spread (estimated using the Stoll-Huang method with price repeats eliminated and the tick test to estimate trade direction) the LIFFE spread measure is slightly smaller than its Eurex counterpart from January, 1997 through April, 1998. Starting in May, 1998, the Eurex traded spread is smaller than the LIFFE traded spread, and the disparity widens through August, 1998, at which time the LIFFE traded spread is more than double the Eurex traded spread.[13]

- The Thompson-Waller spread measure exhibits similar behavior, with small advantages in favor of LIFFE prior to May, 1998, and a widening advantage in favor of Eurex from May, 1998 through the demise of the LIFFE Bund trading in August, 1998.

- The Roll spread measure shows a similar pattern, with Eurex exhibiting a lower spread only beginning in May, 1998. The disparities between LIFFE and Eurex spreads prior to that time are somewhat larger than the disparities across exchanges between the traded and Thompson-Waller spread measures; at times, the Roll spread is more than 20 percent lower on LIFFE.

---

[13]When price repeats (i.e., zero tick trades) are not eliminated, the estimated traded spread on Eurex ranges from .0115 to .0134, and the LIFFE traded spread ranges from .0126 to .0137 prior to June, 1998. The LIFFE spread ballooned to .019 in June, 1998, and to .03 in August, 1998. With price repeats included, from January, 1997-May, 1998, the Eurex spread averaged .0005 less than the LIFFE traded spread. There was no pronounced trend in either the Eurex or LIFFE spreads during the January, 1997-May, 1998 period, nor was there a trend in the Eurex spread post-August, 1998.

- The Eurex spread measures are quite stable, and exhibit no pronounced decline during the period in which LIFFE spread measures widened.

- LIFFE spread measures are also quite stable, exhibiting no pronounced increase during the period of the exchange's market share erosion until the very end of the tipping process, by which time LIFFE market share had fallen below 20 percent.

- Eurex spread measures in the post-August, 1998 period (in which Eurex monopolized the Bund futures trade) are slightly larger than the LIFFE spread measures in the January, 1997-May, 1997 period in which LIFFE had the largest market share in the sample period. This suggests that the open outcry market offers slightly greater maximum potential liquidity than the electronic market.

The results in Table 2 are broadly consistent with those in Table 1. The traded spread (estimated using the Stoll-Huang methodology and the Lee-Ready trade direction estimator) and the total spread (the sum of the realized and effective spreads) fluctuate within a narrow range until the decline in LIFFE volume and market share became pronounced. Indeed, during the April-June, 1998 period, these spreads were slightly smaller than those observed in 1997.

These results suggest that conventional measures of liquidity costs cannot explain the tipping of the Bund market from LIFFE to Eurex. Eurex did not exhibit lower liquidity costs prior to the beginning of the pronounced shift of business from London to Frankfort. Indeed, Eurex liquidity costs exceeded those on LIFFE until the tipping process was almost complete.[14]

This further suggests that other determinants of trading costs–such as exchange fees and brokerage expenses–were important factors in determining the allocation of trading volume. Looking at measured spreads alone, it is difficult to understand how Eurex could have attracted even 35 percent of trading volume; however, Eurex's lower exchange fees of about \$.33 per contract (in 1997) and \$.17 (in 1998) as compared to LIFFE's \$.66 (in 1997) and \$.33 (after March, 1998) (Mehta, 1999) tended to reduce the former exchange's cost disadvantage.

Based on both the traded spread and realized spread estimates, during September-December, 1997 the per contract liquidity cost on Eurex was approximately \$.50 higher on Eurex than LIFFE. However, the January, 1998 price cut eliminated or reversed this difference. Due to the price cut, Eurex's fees were about \$.49 below LIFFE's for traders who spent less than 4500 DM per year in fees, but were \$.66 lower for those who transacted enough to hit the fee cap.[15] The acceleration of Eurex's market share gains coincided with the fee cut.

These differences are very small relative to total trading costs. Over the December, 1997-March, 1998 period, trading cost (spread plus fee) was less than 1 percent for Eurex traders subject to the fee cap. Thus, the market tipped despite the small cost differentials documented here.

Other factors may have exacerbated these cost differentials. Although

[14]Pirrong (1996) and Kofman and Moser (1997) present evidence that differences in liquidity between LIFFE and Eurex were very small as early as 1992.

[15]The fee cuts also sharply reduced the fixed cost of trading on Eurex.

brokerage cost figures are not available, Price-Waterhouse estimated that the cost of contract execution on an open outcry exchange (including labor and equipment costs) was about three times that incurred on an electronic exhange (Mehta, 1999). Given that brokerage fees (as distinct from exchange fees) must cover the cost of contract execution, *ceteris paribus* this cost differential inflated brokerage costs on LIFFE relative to those on Eurex, and contributed further to Eurex's operational cost advantage.[16]

Thus, the data support the inference that out-of-pocket costs (exchange and brokerage fees) were important determinants of the allocation of trading volume across exhanges. Moreover, inasmuch as volume shifted to the exchange with lower non-liquidity related costs, the results are consistent with the implications of the model of Pirrong (2003c). That paper presents a static model in which the lower cost exchange has a competitive advantage, and can undercut its rivals' fees as a result. In equilibrium, the low cost exhange captures 100 percent of the trading volume. In the model, liquidity costs do not depend on which exchange prevails (as is effectively the case here), but total costs are minimized in equilibrium.

4.2 Depth and Price Impact

Panel A of Table 3 presents the average (by month) of the estimated λ coefficients for the Eurex market. These coefficients are estimated using the volume-by-tick data available from Eurex assuming that only trade direction is latent. Panel B of Table 3 presents the average (by month) of the estimated

[16] Kofman and Moser claim that commissions were higher on LIFFE.

λ coefficients for LIFFE and Eurex. These coefficients are estimated from the price data alone assuming that volume is latent. Figure 2 presents a graph illustrating the estimated λ coefficients for LIFFE and Eurex, one observation per week, where these coefficients are estimated assuming volume is latent.

The λ's are somewhat difficult to interpret. Table 4 presents these results in a way that allows a more intuitive interpretation of price impact. Consider a point in time when $\exp(m_t)$, the "true" price, equals 100.05DM, and the market is 100 bid, 100.01 asked. That is, the spread is the minimum tick and the true price is at the quote midpoint. Table 4 reports the trade size (in contracts) required to cause the price to move beyond the current bid or ask, i.e., the trade size that moves price to 100.02 for a buy, or to 99.99 for a sell. This is relevant to traders because smaller trades would occur at the bid or the ask. Thus, the table quantifies market depth at the inside quote when the "true" price is at the quote midpoint. Formally, Table 4 reports:

$$D = \ln(1 + \frac{.5\theta}{100.05})/\lambda$$

where the λ's are those estimated using the latent volume method.

Several results stand out:

- Price impact declined (i.e., depth increased) gradually on Eurex throughout the June, 1997-July, 1998 period. The latent trade direction model and latent volume model exhibit a similar time series pattern of gradual decline throughout 1997 and 1998.

- Price impact remained relatively constant on LIFFE until late 1997, at which time it began to drift upwards. Price impact on LIFFE began to rise dramatically beginning in March, 1998.

- Prior to late-1997, LIFFE price impact was smaller than Eurex price impact. By November-December, 1997 Eurex had achieved rough price impact parity with LIFFE, which continued through February, 1998. In March, 1998 Eurex began to enjoy a substantial depth advantage over LIFFE, which only widened in subsequent months as Eurex depth continued its steady improvement and LIFFE depth eroded noticably.
- Controlling for volume, LIFFE price impact was smaller than Eurex's. This is consistent with the hypothesis that adverse selection costs are higher on an automated trading platform.

The price impact results indicate a substantial deterioration in LIFFE liquidity occurred beginning in March, 1998, somewhat earlier than is evident in the spread measures. Recall that LIFFE spreads began to widen noticably in May at the earliest (for the Roll spread measure) and only somewhat later for other spread measures. In contrast, the depth analysis suggests that LIFFE liquidity began to decline sharply somewhat earlier, when LIFFE's market share was between 30 and 40 percent, and when LIFFE volume had fallen approximately 20 percent relative to its value of May-October, 1997. The decline in depth continued along with the decline in volume.

The different behavior of spreads and depth is consistent with Hasbrouck's (2004) conjecture that these indicators measure different aspects of liquidity, with spreads related mainly to trade processing and inventory costs, and with adverse selection costs driving the depth measure.

Given the estimates of price impact coefficients, it is possible to quantify the effect of depth differences on trading costs. In a market with a discrete tick, finite depth impacts trading costs only for buys (sells) sufficiently large to cause price to move to the next tick above (below) the prevailing ask (bid). Since for even an infinitely deep market a trade occurs at the bid or the ask, finite depth imposes a cost on a liquidity demander only if his trade causes price to jump past the current bid or ask. The size of a trade necessary to cause such a move depends not only on the price impact coefficients, but on the location of the "true" price relative to the bid and ask.

Consider a trade of size $V^* > 0$ (a similar analysis holds for a sale). The current ask is P_A. Define $P_i^*(V^*) = e^{-\lambda_i V^*} P_A$, where λ_i is the price impact coefficient on exchange i (recalling that λ measures the impact of a trade on the log price). That is, $P_i^*(V^*)$ is the level of the true price that would obtain in the absence of a trade such that a purchase of V^* contracts just causes the price to move to the current ask on exchange i. For this trade size, if the no-trade true price exceeds $P_i^*(V^*)$, the trade will cause the transaction price to jump to the next tick level. Defining the true price as $\hat{P}$, the expected cost attributable to finite depth on exchange i is:

$$C_i(V^*) = \int_{P_A-\theta}^{P_A} \theta V^* \Pr(\hat{P} \geq P_i^*(V^*)) d\hat{P}$$

Assuming for simplicity that the true price is uniformly distributed on $[P_A - \theta, P_A]$, this expression becomes:

$$C_i(V^*) = \int_{P_A-\theta}^{P_A} \theta V^* \frac{P_A - P_i^*(V^*)}{\theta} d\hat{P} = \theta V^*(P_A - P_i^*(V^*))$$

The difference in cost between (a) executing a trade of size V^* on Eurex ($i = E$) and (b) executing a trade of the same size on LIFFE ($i = L$) equals $C_E(V^*) - C_L(V^*)$. Assuming a per contract exchange fee difference between LIFFE and Eurex of ΔF, the all-in cost difference equals $C_E(V^*) - C_L(V^*) - \Delta F V^*$. If this expression is negative (positive), it is cheaper to trade on Eurex (LIFFE).

The coefficients reported in Panel B of Table 3 imply that the sum of price impact costs and fees were smaller on LIFFE until November, 1997 for all trade sizes above $V^* = 5$. For bigger trades, the cheaper fees on Eurex did not offset the impact on trading costs of its lower depth. For trades of 100 contracts in October, 1997, the cost advantage on LIFFE was about \$4.00 per contract. However, the substantial relative improvement in Eurex depth in November, 1997 caused the cost advantage to swing in that exchange's favor. Due to the relatively small difference between Eurex and LIFFE price impact coefficients in November, only trades of more than 165 contracts were cheaper to execute on LIFFE. The January, 1998 Eurex price cuts cemented this advantage. During January, traders subject to the fee cap paid \$.50 per contract less in price impact costs and exchange fees on trades of 100 contracts. By March, the relative decline in Eurex price impact meant that its cost advantage was increasing in trade size. For trades of 200 contracts its cost advantage in March, 1998 had ballooned to \$6.00 per contract for traders subject to the fee cap.[17]

---

[17]This analysis, based on the latent volume model, probably overstates the differences in trading costs for large volumes across exchanges, and understates the effect of price cuts on trading cost differentials. Recall that experimentation suggests that the latent volume method overstates price impact coefficients. Thus, the correct $P^*(V^*)$ is bigger for buys and smaller for sales than the value implied by the coefficients estimated using the latent volume method. This in turn implies that the latter coefficients lead to overestimates of the C_i's. Although this could lead to increases in $|C_E - C_L|$, more plausibly it causes this absolute difference to decline. Such a result, in turn, increases the importance of non-liquidity costs in determining all-in trading costs. This analysis also highlights the

The price impact analysis has implications beyond the details of the LIFFE-Eurex struggle. Specifically, it suggests that the relation between volume and price impact/liquidity is quite flat if volume exceeds a threshold level, but becomes quite steep when volume falls below this level. Figure 3 depicts the relation between daily trading volume, and the price impact coefficient estimated for that day. The figure presents results for both exchanges, with Eurex observations denoted by a square and LIFFE observations indicated with a diamond. The figure is based on λ's estimated from the latent volume model.[18] The figure also presents a power law fit of the volume-price impact relation. Note that the relation is very flat for large volumes, but becomes very steep for low volumes. For instance, going from a volume of 120,000 contracts to a volume of 200,000 contracts has only a small effect on price impact (and hence liquidity) but going from a volume of 120,000 to a volume of 40,000 leads to a very large price increase in price impact (and hence a large decline in liquidity). This finding is consistent with the prediction of canonical market microstructure models (see Figure 1 in Pirrong, 2002).

The depth analysis provides some evidence that adverse selection costs are lower, all else equal, in an open outcry trading environment. Since the anonymity of an electronic market makes it easier for the informed to conceal their trading activity than is possible in an open outcry environment, *ceteris*

importance of tick size. A large tick size mitigates the effect of depth differences on the trading cost incurred by liquidity demanders, and as a consequene increases the impact of fee differentials on trading cost differentials.

[18] The relation between Eurex volume and price impact based on the latent trade direction model is quite similar in appearance to Figure 3.

paribus adverse selection should be more severe in the computerized market. In making this comparison, it is important to control for volume as adverse selection costs depend on the level of trading activity. For LIFFE, the average value of the price impact coefficient when daily volume is between 100,000 and 200,000 contracts equals 4×10^{-7}. For Eurex, the average value of this coefficient for this volume range is 1.4×10^{-6}. The LIFFE price impact coefficient is smaller, on average, even when daily volumes exceed 200,000 contracts. Thus, the results suggest that the choice between electronic and open outcry trading involves a trade-off between adverse selection costs and operational efficiencies, with the latter factor favoring the electronic market and the former favoring the open outcry exchange.

5 Summary and Conclusions

This article analyzes a notable example of "tipping" of a market–the shift in Bund futures trading from LIFFE to Eurex during the course of 1997-1998. Prior to the shift in trading volume, liquidity-related trading costs as measured by conventional spread measures were slightly lower on the open outcry LIFFE until that exchange's market share had fallen to less than 20 percent of overall Bund futures trading volume. Only as LIFFE volume fell to effectively nothing did LIFFE liquidity (as quantified by traditional measures of bid-ask spreads) rise substantially absolutely and relative to the comparable costs on the computerized Eurex.

Another measure of liquidity tells a somewhat different story. Specifically, price impact (market depth) coefficients estimated using Markov Chain Monte Carlo methods suggest that the deterioration in LIFFE liquidity became pronounced in March, 1998 and worsened rapidly in the following months. Indeed, consistent with models of financial market liquidity, the sharp fall in LIFFE liquidity (as measured by price impact) accelerated as the exchange's volume loss accelerated in April-June, 1998.

The difference in the behavior of spreads and depth is consistent with Hasbrouck's (2004) suggestion that spread and depth measures of trading costs behave differently because they are related to different aspects of liquidity. In particular, Hasbrouck suggests that the permanent price impact of a transaction reflects adverse selection costs, but that spreads may be driven primarily by trade processing and inventory costs. The empirical findings are consistent with this, as theory suggests that the division of trading volume across trading venues primarily affects adverse selection costs measures. Spreads did not change markedly as tipping proceeded, but adverse selection driven market depth coefficients did change with volume movements as theory predicts. Thus, price impact coefficients are likely the superior measures of the impact of tipping.

The pace of the tipping process–as measured by exchange market share–increased dramatically when the computerized Eurex dramatically cut per trade and fixed fees absolutely and relative to LIFFE's in January, 1998. The biggest single jump in Eurex market share (a 17 point increase) occurred in this month. This suggests that the price cuts were the key factor that accelerated the tipping. The price impact analysis implies that Eurex had achieved a slight advantage in all-in trading cost (on the order of $.15 per contract) on 100 lot trades as early as November, 1997, but that the price cuts trebled this advantage for Eurex users benefitting from the fee cap.

Other evidence is suggestive of the importance of the price cut, but is not sufficiently detailed to permit a more definitive conclusion. Specifically, in the last quarter of 1997 approximately 9 percent of Eurex volume (for all products) originated in the UK. In the first quarter of 1998, this figure rose to 12.84 percent. The percentage was 10.83 percent in January and 12.67 percent in February. Given that (a) volume in futures other than those for German government interest rate products (Bunds and Bobls) was flat during this period, and (b) volumes originating in Germany rose absolutely, it is likely that this increase in the percentage of volume originating in the UK represents a flow of Bund (and Bobl) business to Eurex coincident with the price cuts.

Although the price cuts arguably proved decisive in determining the timing and speed of tipping, they do not explain the steady increase in Eurex volume and market share in 1997, and the corresponding decline in the difference between Eurex and LIFFE price impact coefficients over this period. Without this steady erosion in LIFFE's depth advantage, the January, 1998 price cuts would not have sufficed to make Eurex's all in trading costs lower; based on June, 1997 depth coefficients, for example, even after the price cuts Eurex's trading costs would have been higher for all except relatively small trades. One possible explanation is that the higher brokerage costs associated with floor trading (discussed in section 4 above) induced a gradual diversion of some trading activity to Eurex. In addition, US-orgin order flow to Eurex increased due to the placement of Eurex terminals in the US beginning in March, 1997; this reduced the cost of accessing the Eurex market from the

US absolutely and relative to the cost of accessing LIFFE.[19] Together, these additional volumes reduced LIFFE's liquidity advantage sufficiently to make the market ripe for tipping via aggressive price cuts.

This analysis draws attention to the importance of non-liquidity related costs in determining the victorious trading venue. Academics tend to focus on liquidity as the key element of trading cost. The data from LIFFE and Eurex suggest that liquidity costs are the largest single component of trading costs, but they also suggest that liquidity cost differences across trading venues may be quite small when each exchange has substantial trading volume. Under these circumstances, small differences in non-liquidity-related costs can be decisive. This is consistent with the predictions of Pirrong (2002, 2003a-c), and with the view expressed by a former Chairman of the Chicago Mercantile Exchange:

> Critical mass goes this way. The competitive market has 20 percent market share, that's okay. When it has 30 percent, one starts to take notice . . . . At 45 percent you've already lost. *The reason it has moved from 30 to 45 percent is because of cost efficiencies being recognized.* . . . Before you know it, in no time

[19] By October, 1997 18 percent of Eurex Bund business originated in the US (Young and Theys, 1999). This represents about 33,000 contracts per day. Of course, US FCMs might have directed some of this volume to Eurex even if trading terminals had not been approved for the US, so it represents an upper bound on the additional volume going to Eurex due to the installation of trading terminals there. However, it is interesting to note that US share of total Eurex volume across all contracts was negligible (.1 percent) in January, 1997 (prior to the installation of terminals in the US), but had grown to 2.3 percent by October, and to around 5 percent in the first quarter of 1998.

it goes to 70 percent. Game over.[20]

That is, operational efficiencies (reflected in lower fees and/or lower brokerage costs) can induce market tipping when the more operationally efficient exchange's volume approaches, but is still smaller than, that of the putatively less operationally efficient incumbent.

This begs the question of how the putatively more efficient entrant can generate sufficient trading volume so that its operating cost advantages can off-set higher liquidity costs. Note that the power law relation between volume and price impact implies that an exchange with a small volume faces a substantial liquidity cost disadvantage relative to the larger exchange.

The Eurex experience provides one answer to this question–strong patrons who direct considerable order flow to the exchange. The statistical analysis presented herein does not start at the beginning, but at the beginning of the end. Eurex had grown prior to 1997 due primarily to the patronage of German banks that just happened to be the owners of the exchange.[21] These banks provided sufficient volume to permit Eurex to survive and provide liquidity-related trading costs that were at worst only slightly higher than those on LIFFE. This put Eurex within striking distance of LIFFE. Lower access costs due to the operational efficiencies of an electronic market appar-

[20]Emphasis added. Quote from Barret and Scott, who interviewed numerous exchange representatives in London and Chicago in 1998 and 1999. The quote is anonymous, but is attributed to an "ex-Chairman CME," and is therefore most likely Leo Melamed, Jack Sander, or Brian Monieson.

[21]In 1997, 80.8 percent of Eurex volume was traded by German members, the largest of which were the German banks that owned Eurex. By the end of 1998, German members accounted for only about 50 percent of volume; by mid-1999 this fell to about 40 percent.

ently attracted additional business (especially from the United States), which narrowed the liquidity cost differential further throughout the course of 1997. When the liquidity differential was nearly closed by the end of that year, Eurex was positioned to deliver the *coup de grâce* with its fee cut–assisted not a little by LIFFE's delayed response.

Absent support like the German banks provided to Eurex, it is by no means clear how a more operationally efficient entrant exchange (that is, an entrant that would, if all trading activity took place there, offer lower total trading costs) can generate the volume necessary for liquidity cost differences to fall below the entrant's operating cost advantage. This is consistent with Pirrong (2003c), who argues that the structure of the industry that exerts control over order flow can be the decisive factor in determining the victor in a winner-take-all competition between exchanges. The existence of large brokerage or trading firms that control large order flows, and which can coordinate their activities, increases the likelihood that an entrant can wrest a market from a well-established incumbent. The German banks played this role in their support of DTB/Eurex.

This is also consistent with the importance accorded technology sponsorship in the literature on network industries. The German banks effectively sponsored the Eurex platform. Without this sponsorship, it is highly doubtful whether Eurex would have achieved the critical mass necessary to permit it to reach rough parity in liquidity cost with LIFFE. Once Eurex's liquidity costs were only slightly higher than LIFFE's, it could exploit its operational efficiencies through price cutting.

Eurex's experience in its attempt to compete with the Chicago Board of

Trade for dominance in US Treasury bond and note futures provides further evidence of the importance of pricing and order flow control. In 2003 Eurex entered the US market. It employed a penetration pricing strategy focused on attracting large order flows. This pricing strategy promised large fee rebates to the 10 largest customers. The tournament-like payoff structure was designed to give disproportionately large fee cuts to the brokerage firms that provided the largest volumes to the exchange, and was specifically designed to attract big blocks of orders.

So far, Eurex has made very little headway in its attempt to enter the US market in large part because the Chicago Board of Trade apparently absorbed the lessons of LIFFE's near death experience. Rather than repeating LIFFE's mistake of delaying cutting its trading fees in response to Eurex price competition, the CBOT preemptively cut trading fees immediately prior to Eurex's launch of Treasury futures trading. Moreover, in contrast to its experience with the Bund where exchange owners controlled large order flows, in the US Eurex did not possess the core of loyal order flow suppliers that had been key to its success in Bunds. These two factors–the CBOT's aggressive pricing strategy and the lack of "sponsorship"–help explain why Eurex has attracted only a trivial share of Treasury futures trading during its first months of operation in the US.

As another example of attempted inter-exchange competition, EuronextLIFFE has offered dollar denominated short term interest rate contracts in competition with the incumbent Chicago Mercantile Exchange. Although EuronextLIFFE has fared somewhat better in the US than Eurex, the CME's volume erosion has still been quite small. CME has also been quite aggressive in setting trading fees, and EuronextLIFFE did not enter the competition with a strong cadre of order flow suppliers.

Together, the Eurex-LIFFE, EUREX-CBOT, and Euronext-CME experiences suggest that an entrant exchange can wrest a futures contract from an incumbent only under some rather special circumstances–specifically, the entrant must control sufficient order flow to make its liquidity costs comparable to the incumbent's, and the incumbent must not respond to the entrant's price cutting. The first condition is very expensive to achieve, and the second requires a strategic error on the part of the incumbent. Thus, tipping can be expected to be the exception, rather than the rule.

Beyond its implications for the nature of competition between exchanges, this article is an interesting case study of competition in a network industry. Although putative examples of network effects and tipping are legion (especially in the computer and telecommunications industries), it is seldom possible to quantify the relative costs of competing networks and to observe how these costs evolve over time in response to consumer choices. This is not the case in futures trading–given the requisite high frequency data and appropriate econometric techniques, it is possible to quantify the impact of network effects. The Eurex-LIFFE episode also demonstrates the importance of network sponsorship and how seemingly trivial price differences can have a decisive impact on who prevails in a contest between two similarly sized networks.

A Introduction

This appendix sets out the MCMC methodology applicable when volume is latent.[22] Although the reported results are for the model that takes price discreteness (tick size) into account, I first describe the methodology that ignores price discreteness because this provides a better understanding of the basics of the estimation approach.

B Estimation Without Discreteness

Estimation of β and λ involves a straightforward implementation of a Gibbs Sampler when transactions can occur at any price on the positive real line. The procedure is as follows.

Remark. A glance at (3) indicates that without additional information, the model is not identified. Multiplying each V_t by a constant k and dividing both β and λ by k would produce the same value of the likelihood function. I defer discussion of this issue until Appendix D. In brief, I utilize daily volume data (which is available) to identify the model.

λ, β, σ^2, $\{m_t\}$ and $\{V_t\}$ are parameters in the model. Parameters are "blocked," with λ, β, and σ^2 in one block, $\{V_t\}$ in a second, $\{m_t\}$ in a third.

I assume that the prior distributions of β and λ are normal, with both parameters constrained to be non-negative. The non-negativity constraint is economically plausible, and is necessary to identify the model (as noted

[22]See Hasbrouck (2002) for estimation when only trade direction is latent, but (unsigned) volume is available. Although Hasbrouck does not account for heteroskedasticity attributable to variable time between trades, I modify his methodology to do so.

in Hasbrouck, 2002). As is conventional with a normal ε_t, I assume an inverted gamma distribution as the prior for σ^2. With the exception of the non-negativity constraint, the priors are non-informative.

The first step of the procedure is to choose $\{V_t\}$ and $\{m_t\}$ (arbitrarily, though the choice of the m_t is consistent with the observed p_t and V_t). This choice, the priors, and standard Bayesian results for linear regressions applied to (3) imply posterior distributions for λ, β, and σ^2. Since the errors are heteroskedastic, but the variance-covariance matrix is the product of a known matrix and σ^2, the posterior can be estimated using expressions by multiplying the both sides of (3) by the matrix A, where

$$A = \begin{pmatrix} 1/\sqrt{\tau_1} & 0 & \cdots & \cdots & \cdots \\ 0 & 1/\sqrt{\tau_2} & 0 & \cdots & \cdots \\ 0 & 0 & 1/\sqrt{\tau_3} & 0 & \cdots \\ \vdots & \vdots & \vdots & \ddots & \cdots \\ 0 & 0 & 0 & \cdots & 1/\sqrt{\tau_T} \end{pmatrix}$$

Draws are then made from these posteriors.

Given the newly drawn regression parameters, the Gibbs sampler next takes draws for $\{V_t\}$ and $\{m_t\}$. Values are drawn sequentially for each t.

First, draw $V_t|m_{t-1}, m_{t+1}, V_{t-1}, V_{t+1}, p_t$. By Bayes' Rule:

$$Pr(V_t|m_{t-1}, m_{t+1}, V_{t+1}, p_t) \propto Pr(p_t|m_{t-1}, m_{t+1}, V_t, V_{t+1})Pr(V_t|m_{t-1}, m_{t+1}, V_{t+1})$$

Consider the second term on the RHS. First note:

$$m_t = m_{t-1} + \lambda V_t + \varepsilon_t\sqrt{\tau_t}$$

$$m_{t+1} = m_t + \lambda V_{t+1} + \varepsilon_{t+1}\sqrt{\tau_{t+1}}$$

Therefore:

$$V_t = \frac{1}{\lambda}[m_{t+1} - m_{t-1} - \lambda V_{t+1} - \varepsilon_t\sqrt{\tau_t} - \varepsilon_{t+1}\sqrt{\tau_{t+1}}] \tag{4}$$

This implies:

$$V_t \sim N(\frac{1}{\lambda}[m_{t+1} - m_{t-1} - \lambda V_{t+1}], \frac{\sigma^2}{\lambda^2}(\tau_t + \tau_{t+1}))$$

Now consider the first term.

$$p_t = m_{t-1} + \lambda V_t + \beta V_t + \varepsilon_t \sqrt{\tau_t}$$

Also,

$$m_{t+1} = m_t + \lambda V_{t+1} + \varepsilon_{t+1} \sqrt{\tau_{t+1}}$$

Therefore, since $p_t = m_t + \beta V_t$:

$$p_t = m_{t+1} - \lambda V_{t+1} + \beta V_t - \varepsilon_{t+1} \sqrt{\tau_{t+1}}$$

Averaging the two previous expressions for p_t produces:

$$p_t = .5[m_{t+1} + m_{t-1} + \lambda(V_t - V_{t+1}) + 2\beta V_t + \varepsilon_t \sqrt{\tau_t} - \varepsilon_{t+1} \sqrt{\tau_{t+1}}] \quad (5)$$

Consequently,

$$p_t | m_{t-1}, m_{t+1}, V_t, V_{t+1} \sim N(.5[m_{t+1} + m_{t-1} + \lambda(V_t - V_{t+1}) + 2\beta V_t], \frac{\sigma^2}{4}(\tau_t + \tau_{t+1}))$$

Thus, $Pr(V_t | m_{t-1}, m_{t+1}, V_{t+1}, p_t)$ is a convolution of normals, from which it is straightforward to draw V_t.

One way to do this is to rewrite (5) as follows:

$$V_t = \frac{2p_t - m_{t+1} - m_{t-1} + \lambda V_{t+1} - \varepsilon_t \sqrt{\tau_t} + \varepsilon_{t+1} \sqrt{\tau_{t+1}}}{\lambda + 2\beta}$$

and combine this with (4) to obtain:

$$V_t = \frac{1}{\lambda(\lambda + 2\beta)}[\lambda p_t + \beta m_{t+1} - (\lambda + \beta) m_{t-1} - \lambda\beta V_{t+1} - (\lambda + \beta)\varepsilon_t \sqrt{\tau_t} - \beta \varepsilon_{t+1} \sqrt{\tau_{t+1}}]$$

This implies:

$$V_t \sim N(\hat{V}_t, \sigma_V^2)$$

where

$$\hat{V}_t = \frac{[\lambda p_t + \beta m_{t+1} - (\lambda + \beta) m_{t-1} - \lambda\beta V_{t+1}]}{\lambda(\lambda + 2\beta)}$$

and

$$\sigma_V^2 = \sigma^2 \frac{\tau_t(\lambda + \beta)^2 + \tau_{t+1}\beta^2}{\lambda^2(\lambda + 2\beta)^2}$$

In the absence of discreteness, (2) implies that since p_t is known, given V_t and β,

$$m_t = p_t - \beta V_t$$

Once V_t and m_t have been chosen, it is possible to determine new posteriors for λ, β, and σ^2 *via* (1). After drawing from these posteriors, one draws the V_t and m_t. This process is repeated S times. After discarding the first S_1 draws of the parameters, the remaining $S - S_1$ draws are used to draw inferences about the parameters.

C Estimation With Discreteness

Discreteness of the price implies that transaction prices for purchases ($V_t > 0$) are rounded up to the nearest price increment, and prices for sales ($V_t < 0$) are rounded down to the nearest price increment. Discreteness requires two modifications of the methodology, and makes it necessary to condition on the dollar price $P_t = exp(p_t)$ rather than on the log price p_t.

The first modification relates to the draw of $V_t | m_{t-1}, m_{t+1}, V_{t+1}, P_t$. Specifically, discreteness alters $Pr(P_t | m_{t-1}, m_{t+1}, V_t, V_{t+1})$. To derive this density, denote θ as the tick size, and define

$$p_t^* = m_t + \beta V_t.$$

This is the unobserved "true" log price that would be observed in the absence of discreteness. For $V_t > 0$, if a trade takes place at P_t, then $ln(P_t - \theta) \leq p_t^* \leq p_t$ because for buys price is rounded up to the nearest tick. For $V_t < 0$, price is rounded down to the nearest tick, implying $p_t \leq p_t^* \leq ln(P_t + \theta)$. Therefore, for $V_t > 0$, $ln(P_t - \theta) - \beta V_t \leq m_t \leq p_t - \beta V_t$, and for $V_t < 0$, $p_t - \beta V_t \leq m_t \leq ln(P_t + \theta) - \beta V_t$. Therefore, for $V_t > 0$,

$$Pr(P_t|m_{t-1}, m_{t+1}, V_t, V_{t+1}) = Pr(m_t \in [ln(P_t - \theta) - \beta V_t, p_t - \beta V_t]|m_{t-1}, m_{t+1}, V_t, V_{t+1})$$

while for $V_t < 0$

$$Pr(P_t|m_{t-1}, m_{t+1}, V_t, V_{t+1}) = Pr(m_t \in [p_t - \beta V_t, ln(P_t + \theta) - \beta V_t]|m_{t-1}, m_{t+1}, V_t, V_{t+1})$$

Note that

$$m_t|m_{t-1}, m_{t+1}, V_t, V_{t+1} \sim N(\hat{m}_t, \frac{\sigma^2}{4}(\tau_t + \tau_{t+1}))$$

where

$$\hat{m}_t = .5[m_{t+1} + m_{t-1} + \lambda(V_t - V_{t+1}) + 2\beta V_t]$$

Therefore, for $V_t > 0$

$$Pr(P_t|m_{t-1}, m_{t+1}, V_t, V_{t+1}) = \int_{\underline{m}^+}^{\overline{m}^+} n(m, \hat{m}_t, \frac{\sigma^2}{4}(\tau_t + \tau_{t+1}))dm$$

where $n(a, b, c)$ indicates the density of a normal variate with mean b and variance c evaluated at a, $\underline{m}^+ = ln(P_t - \theta) - \beta V_t$ and $\overline{m}^+ = p_t - \beta V_t$.

For $V_t < 0$,

$$Pr(P_t|m_{t-1}, m_{t+1}, V_t, V_{t+1}) = \int_{\underline{m}^-}^{\overline{m}^-} n(m, \hat{m}_t, \frac{\sigma^2}{4}(\tau_t + \tau_{t+1}))dm$$

where $\underline{m}^{-} = p_t - \beta V_t$ and $\overline{m}^{-} = ln(P_t + \theta) - \beta V_t$.

These integrals can be rewritten as follows:

$$
\begin{aligned}
Pr(P_t|m_{t-1}, m_{t+1}, V_t, V_{t+1}) &= N(\overline{m}, \hat{m}_t, \frac{\sigma^2}{4}(\tau_t + \tau_{t+1})) \\
&- N(\underline{m}, \hat{m}_t, \frac{\sigma^2}{4}(\tau_t + \tau_{t+1}))
\end{aligned}
$$

where $N(a, b, c)$ is the cumulative normal distribution with mean b and variance c evaluated at a, $\overline{m} = \overline{m}^{+}$ and $\underline{m} = \underline{m}^{+}$ for $V_t > 0$, and $\overline{m} = \overline{m}^{-}$ and $\underline{m} = \underline{m}^{-}$ for $V_t < 0$. Note that $\overline{m}$ and $\underline{m}$ both depend on V_t.

This analysis implies that $Pr(V_t|m_{t-1}, m_{t+1}, V_{t-1}, V_{t+1}, P_t)$ is proportional to the product of a normal density and the difference between two cumulative normals. This is a non-standard density that cannot be drawn from directly. However, it is readily drawn from using a Griddy-Gibbs approximation to the density. Alternatively, realistically constraining volumes to take integer values facilitates drawing from the posterior. The density of interest, $Pr(V_t|P_t, m_{t-1}, m_{t+1}, V_t, V_{t+1})$, is proportional to the product of a normal density and the the difference between cumulative normals. To estimate this density, I calculate the product for each integer value of V_t, with $|V_t| \in [1, \bar{V}]$, and normalize the product for each integer value by their sum. This gives a discrete density for volume. The vector of possible integer values of V_t is readily created. This vector can be used to create $\hat{m}_t$, $\underline{m}$, and $\overline{m}$. Therefore, it is straightforward to calculate the relevant products for each V_t value. These can be summed to normalize. The main computational cost is calculating the cumulative probability, which requires looping through the entire probability vector and taking the cumulative sum at each step in the loop.

Given the draw for V_t, m_t is drawn from $N(\hat{m}_t, \frac{\sigma^2}{4}(\tau_t + \tau_{t+1}))\mathbf{1}_{\mathbf{m}\in[\underline{\mathbf{m}},\overline{\mathbf{m}}]}$. Given the $\{m_t\}$ and V_t, the posteriors for λ and σ^2 are determined by the regression:

$$\Delta m_t = \lambda V_t + \varepsilon_t \sqrt{\tau_t}$$

The second modification of the methodology required to address price discreteness relates to the draw of β. Note that the blocking is slightly different with discreteness. In the absence of price discreteness, it is straightforward to draw from the posterior of λ and β implied by the regression of price changes on volume and the change in volume. This is not possible with discrete prices, because p_t^* is not observed. This requires a different parameter blocking scheme and the employment of a Metropolis-Hastings algorithm.

With respect to blocking, whereas without discreteness, β, λ, and σ^2 are blocked together, with discreteness λ and σ^2 are blocked together, and β is in a separate block.

Metropolis-Hastings requires the choice of a candidate density for β. I employ a Random Walk Markov Chain method to make initial draws of β. Specifically, if $\beta^{(i)}$ is the value of β in iteration i of the sampler, the candidate for $\beta^{(i)}$ is drawn from:

$$\beta^{(i)} = \beta^{(i-1)} + \mu_t$$

where μ_t is a normal variate with standard deviation σ_μ, and where $\beta^{(i)}$ is constrained to be non-negative. Denote this density as $G(\beta)$. In essence, this method randomly chooses a new value for β. The Metropolis-Hastings algorithm is used to determine whether to use the new $\beta^{(i)}$ or to instead discard the new draw and continue to use $\beta^{(i-1)}$. The MH algorithm accepts the new draw randomly, and the new draw is more likely to be accepted if it results in a higher likelihood than the old value $\beta^{(i-1)}$. Thus, the method wanders in the β dimension, but tends to concentrate its time in areas with high values of the likelihood function.

As in Hasbrouck (2002), with a new choice of β, the already selected values of $\{m_t\}$ may violate the bounds implied by discreteness and the observed prices. I therefore adjust the $\{m_t\}$ deterministically to ensure that they do not violate these bounds. Thus, this step (as in Hasbrouck) essentially involves a joint draw of $\{m_t\}$ and β. The new draw of β at step i of the Gibbs sampler is denoted by β^*, and the new $\{m_t\}$ by $\{m_t^*\}$. The draws at step $i-1$ of the sampler are denoted by $\beta^{(i-1)}$ and $\{m_t^{(i-1)}\}$.[23]

Draws from the candidate density at iteration i of the Gibbs sampler are accepted with probability:

$$\alpha = \min[1, \frac{G(\beta^*)Pr(\{m_t^*\}|\{V_t\},\sigma^2,\lambda)f(\beta^*)}{G(\beta^{(i-1)})Pr(\{m_t^{(i-1)}\}|\{V_t\},\sigma^2,\lambda)f(\beta^{(i)})}]$$

where $f(\beta)$ is the prior distribution for β.

In addition to depending on the ratio of the candidate densities at the proposed and existing draws, the acceptance probability depends on the ratio of the posterior probability for β and $\{m_t\}$. Note:

$$Pr(\beta,\{m_t\}|\sigma^2,\{V_t\},P_t,\lambda) = \frac{Pr(P_t|\{m_t\},\{V_t\},\beta,\sigma^2,\lambda)Pr(\{m_t\}|\{V_t\},\beta,\sigma^2,\lambda)Pr(\{V_t\}|\beta,\sigma^2,\lambda)Pr(\beta,\lambda,\sigma^2)}{Pr(P_t,\{V_t\}|\sigma^2,\lambda)Pr(\sigma^2,\lambda)}$$

Note further that the relation between P_t and $\{m_t\},\{V_t\}$, λ and β is deterministic, hence the first term in the numerator equals 1. Due to the

[23] The adjustment in $\{m_t\}$ works as follows. Take the just-drawn m_t, and define the new $m_t^* = m_t + \beta^{(i-1)}V_t - \beta^* V_t$.

independence of the priors,

$$Pr(\{V_t\}|\beta,\sigma,\lambda)Pr(\beta,\lambda,\sigma^2) = Pr(\{V_t\})Pr(\beta)Pr(\sigma^2,\lambda)$$

Consequently,

$$Pr(\beta,\{m_t\}|\sigma^2,\{V_t\},P_t,\lambda) \propto Pr(\{m_t\}|\{V_t\},\beta,\sigma^2,\lambda)f(\beta) = Pr(\{m_t\}|\{V_t\},\sigma^2,\lambda)f(\beta)$$

where the last equality follows from the fact that the distribution of $\{m_t\}$ does not depend on β. Thus, the Pr terms that appear in the acceptance probability formula are proportional to the posterior probability of observing β and $\{m_t\}$. The acceptance probability is therefore the product of the ratios of the candidate densities at the proposed and existing draws and the ratio of the posterior probabilities for the proposed and existing draws.

D Identification

As noted above, the model is not identified. What is needed is some way of tying the estimated $\{V_t\}$ to actual volume levels to ensure that λ and β are scaled properly to reflect the price impact of actual-sized trades.

Although futures markets do not typically report volume trade-by-trade in the high frequency data that they distribute, they almost always report daily volumes. Thus, I tie $\{V_t\}$ draws to actual volume levels by rescaling the V_t draws so that the sum of their abolute values on each day sum to the observed volume on that day. Thus, if $\mathcal{T}_d$ is the set of trades that take place on day d, and $\overline{V}_d$ is the total volume observed on that day, define:

$$k_d = \frac{\sum_{t\in\mathcal{T}_d}|V_t|}{\overline{V}_d}$$

and then set $V_t = V_t/k_d$, and use these volumes in the estimation of β and λ, and in the subsequent draws for the volumes.

Extensive experimentation suggests that this method of identification works effectively. I have estimated the latent trade direction and latent volume models for several markets for which trade-by-trade volume data is available. These markets include not only Eurex, but several Chicago Mercantile Exchange equity index, currency, and agricultural futures contracts. A comparison of the results across estimation methods indicates that the latent volume method tends to generate larger trade impact coefficients than the latent trade direction method, and that the volume draws generated by the latent volume method exhibit fewer very large trades than occur in practice. This finding obtains even when one allows for trade impacts that are concave in trade size (by positing a buy-sell dummy or using the square root of volume as the trade size variable as in Hasbrouck or both). Nonetheless, the two methods generally give the same ranking of trade impact across markets and over time, and the λ coefficients estimated using the two methods are highly and positively correlated. For instance, for the Eurex Bund futures data considered here, the correlation between the latent volume λ and the latent trade direction λ is .7.

Table 1 Panel A Eurex Spread Measures						
Contract	Month	N	Traded	γ	TW	Roll
H97	Jan	35759	.0223	.9736	.0113	.0097
H97	Feb	23434	.0215	.9814	.0109	.0100
M97	Mar	28792	.0223	.9774	.0113	.0096
M97	Apr	33791	.0214	.9748	.0109	.0099
M97	May	26272	.0216	.9769	.0110	.0097
U97	Jun	31126	.0215	.9789	.0109	.0099
U97	Jul	37382	.0209	.9758	.0107	.0104
U97	Aug	36897	.0216	.9757	.0109	.0098
Z97	Sep	28568	.0212	.9868	.0107	.0105
Z97	Oct	48782	.0215	.9750	.0110	.0101
Z97	Nov	24086	.0213	.9868	.0108	.0105
H98	Dec	15807	.0214	.9933	.0108	.0098
H98	Jan	37997	.0214	.9955	.0107	.0103
H98	Feb	37987	.0212	.9883	.0107	.0113
M98	Mar	27892	.0223	.9774	.0113	.0096
M98	Apr	48473	.0214	.9943	.0108	.0120
M98	May	37200	.0211	.9948	.0106	.0120
U98	Jun	39376	.0212	.9966	.0107	.0121
U98	Jul	43597	.0214	.9989	.0107	.0126
U98	Aug	70205	.0222	.9893	.0111	.0124
Z98	Sep	94832	.0218	.9922	.0111	.0100
Z98	Oct	129416	.0223	.9874	.0116	.0108
Z98	Nov	66497	.0215	.9934	.0109	.0101
H99	Dec	36015	.0220	.9932	.0112	.0099
H99	Jan	84838	.0215	.9884	.0110	.0103
H99	Feb	99143	.0216	.9941	.0109	.0102

Note: This panel reports spread measures for Eurex based on the tick test for trade direction identification. “Contract” is the contract month (contract month code and year). “Month” is the calendar month in which the spread is estimated. “N” is the number of trades included (only trades at non-zero ticks included). "Traded" is the Stoll traded spread. "γ" is the fraction of the traded spread attributable to adverse selection and inventory costs. "TW" is the Thompson-Waller spread. "Roll" is the Roll spread.

Table 1 Panel B LIFFE Spread Measures						
Contract	Month	N	Traded	γ	TW	Roll
H97	Jan	42230	.0211	.9977	.0105	.0091
H97	Feb	30948	.0207	.9989	.0104	.0091
M97	Mar	31585	.0211	.9999	.0105	.0080
M97	Apr	35052	.0207	.9945	.0104	.0082
M97	May	31460	.0209	.9973	.0105	.0079
U97	Jun	31197	.0210	.9998	.0105	.0086
U97	Jul	36897	.0206	.9998	.0103	.0085
U97	Aug	31472	.0201	.9997	.0104	.0076
Z97	Sep	25501	.0209	.9982	.0104	.0076
Z97	Oct	42161	.0210	.9950	.0105	.0099
Z97	Nov	23299	.0206	.9986	.0103	.0088
H98	Dec	16387	.0214	.9988	.0107	.0093
H98	Jan	37573	.0209	.9962	.0104	.0086
H98	Feb	37920	.0205	.9927	.0103	.0096
M98	Mar	15443	.0215	.9889	.0108	.0094
M98	Apr	18673	.0211	.9982	.0106	.0073
M98	May	15443	.0215	.9889	.0108	.0138
U98	Jun	8118	.0240	.9985	.0121	.0122
U98	Jul	3707	.0275	.9703	.0138	.0149
U98	Aug	3319	.0494	.9824	.0274	.0133

Note: This panel reports spread measures for LIFFE based on the tick test for trade direction identification. “Contract” is the contract month (contract month code and year). “Month” is the calendar month in which the spread is estimated. “N” is the number of trades included (only trades at non-zero ticks included). “Traded” is the Stoll traded spread. “γ” is the fraction of the traded spread attributable to adverse selection and inventory costs. “TW” is the Thompson-Waller spread. “Roll” is the Roll spread.

Table 2 LIFFE Spread Measures						
Contract	Month	N	Traded	λ	TW	Roll
H97	Jan	16250	.0138	.8876	.0108	.0052
H97	Feb	13132	.0139	.8926	.0109	.0051
M97	Mar	12088	.0139	.9306	.0105	.0041
M97	Apr	13419	.0138	.9037	.0109	.0057
M97	May	12126	.0139	.8877	.0110	.0027
U97	Jun	12333	.0134	.9077	.0108	.0054
U97	Jul	13194	.0139	.9410	.0107	.0043
U97	Aug	12816	.0141	.9646	.0107	.0037
Z97	Sep	9592	.0128	.9005	.0104	.0052
Z97	Oct	5291	.0142	.9286	.0105	.0045
Z97	Nov	8336	.0127	.8485	.0107	.0058
H98	Dec	5291	.0128	.8930	.0107	.0035
H98	Jan	14865	.0136	.9755	.0107	.0046
H98	Feb	11096	.0135	.9755	.0107	.0039
M98	Mar	10309	.0135	.8932	.0107	.0038
M98	Apr	6959	.0129	.9227	.0106	.0042
M98	May	4894	.0127	.9363	.0107	.0041
U98	Jun	1769	.0106	.8929	.0109	.0015
U98	Jul	473	.0137	.9812	.0109	.0063
U98	Aug	694	.0259	.9998	.0120	.0011

Note: This table reports the spread measures for LIFFE based on bid-ask quotes. "Contract" is the contract month (contract month code and year). "Month" is the calendar month in which the spread is estimated. "N" is the number of trades included (only trades at non-zero ticks included). "Traded" is the Stoll traded spread. "γ" is the fraction of the traded spread attributable to adverse selection and inventory costs. "TW" is the Thompson-Waller spread. "Roll" is the Roll spread.

Table 3 Panel A Price Impact Coefficients Latent Trade Direction Method	
Month	EUREX
Jun 97	2.09-E07
Jul 97	1.23E-07
Aug 97	1.34E-07
Sep 97	8.86E-08
Oct 97	1.83E-07
Nov 97	1.35E-07
Dec 97	3.54E-07
Jan 98	7.91E-08
Feb 98	1.04E-07
Mar 98	6.08E-08
Apr 98	7.26E-08
May 98	7.14E-08
Jun 98	8.60E-08

Table 3 Panel B Price Impact Coefficients Latent Volume Method		
Month	EUREX	LIFFE
Jun 97	5.14-E07	9.31E-08
Jul 97	4.62E-07	9.68E-08
Aug 97	5.97E-07	1.28E-07
Sep 97	4.32E-07	1.34E-07
Oct 97	4.61E-7	1.43E-07
Nov 97	2.45E-07	2.31E-07
Dec 97	2.24E-07	3.86E-07
Jan 98	2.34E-07	2.03E-07
Feb 98	1.78E-07	2.21E-07
Mar 98	1.18E-07	3.07E-07
Apr 98	1.78E-07	6.98E-07
May 98	7.64E-08	7.21E-07
Jun 98	9.25E-08	1.44E-06

Table 4 Number of Contracts to Move Price One-Half Tick		
Month	EUREX	LIFFE
Jun 97	97.321	536.883
Aug 97	108.176	516.340
Sep 97	83.750	390.018
Oct 97	115.872	373.592
Nov 97	108.457	349.153
Dec 97	203.972	216.679
Jan 98	223.528	129.672
Feb 98	213.237	246.633
Mar 98	280.498	226.752
Apr 98	424.438	162.876
May 98	280.168	71.673
Jun 98	654.648	69.395
Jul 98	540.385	34.616

References

[1] Admati, Anat, and Paul Pfleiderer. 1988. "A Theory of Intraday Patterns: Volume and Price Variability," *Review of Financial Studies* 3-40.

[2] Admati, A., and P. Pfleiderer. "Sunshine Trading and Financial Market Equilibrium." *Review of Financial Studies* 4 (1991): 443-481.

[3] Breen, W., L. Glosten, and R. Jaganathan. "The Economic Significance of Predictable Variations in Stock Index Returns." *Journal of Finance* 44 (1989): 1177-1189.

[4] Chan, K. C., K. Chan, and G Karolyi. "Intraday Volatility in the Stock Index and Stock Index Futures Markets." *The Review of Financial Studies* 4 (1991): 657-684.

Glosten, L. "Is the Electronic Limit Order Book Inevitable?" *Journal of Finance* 49 (1994): 1127-1161.

[5] Grossman, S. "The Information Role of Upstairs and Downstairs Markets." *Journal of Business* 65 (1992) 509-529.

[6] Hasbrouck, J. "Markov Chain Methods for Bayesian Estimation of Market Microstructure Models." NYU working paper (2002).

[7] Hasbrouck, J. "Trading Costs and Returns for US Equities: Evidence from the Daily Data." NYU working paper (2004).

[8] Huang, R., and H. Stoll. "The Components of the Bid-Ask Spread: A General Approach." *Review of Financial Studies* 10 (1997): 995-1034.

[9] Jones, C., G. Kaul, and M. Lipson. "Transactions, Volume, and Volatility." *Review of Financial Studies* 7 (1994): 631-651.

[10] Kofman, P., and J. Moser. "Spreads, Information Flows and Transparency Across Trading Systems." *Applied financial Economics* 7 (1977): 281-294.

[11] Kyle, A. "Continuous Auctions and Insider Trading." *Econometrica* 53 (1985): 317-355.

[12] Madhavan, A. "Market Microstructure: A Survey." *Journal of Financial Markets* 3 (2000): 205-258.

[13] Mehta, N. "Meridien on the Electronic Threat." http://www.derivativestrategy.com/magazine/archive/1999/0399shrt.asp (1999).

[14] Pagano, M. "Trading Volume and Asset Liquidity." *Quarterly Journal of Economics* 104 (1989): 255-274.

[15] Pirrong, C. "Market Liquidity and Depth on Computerized and Open Outcry Trading Systems: A Comparison of DTB and LIFFE Bund Contracts." *Journal of Futures Markets* 16 (1996): 519-543.

[16] Pirrong, C. "The Industrial Organization of Financial Markets" Theory and Evidence." *Journal of Financial Markets* 2 (1999): 329-358.

[17] Pirrong, C. "A Theory of Financial Exchange Organization." *Journal of Law and Economics* 43 (2000): 437-472.

[18] Pirrong, C. "Security Market Macrostructure." *Journal of Law, Economics, and Organization* 18 (2002): 385-410.

[19] Pirrong, C. "The Macrostructure of Electronic Financial Markets." Working paper, University of Houston (2003a).

[20] Pirrong, C. "Third Markets and the Second Best." Working paper, University of Houston (2003b).

[21] Pirrong, C. "Upstairs, Downstairs: Electronic vs. Open Outcry Exchanges." Working paper, University of Houston (2003c).

[22] Young, P., and T. Theys. *The Capital Market Revolution: The Future of Markets in an Online World.* New York: Financial Times-Prentice Hall (1999).

Figure 1
DTB and LIFFE Bund Market Shares

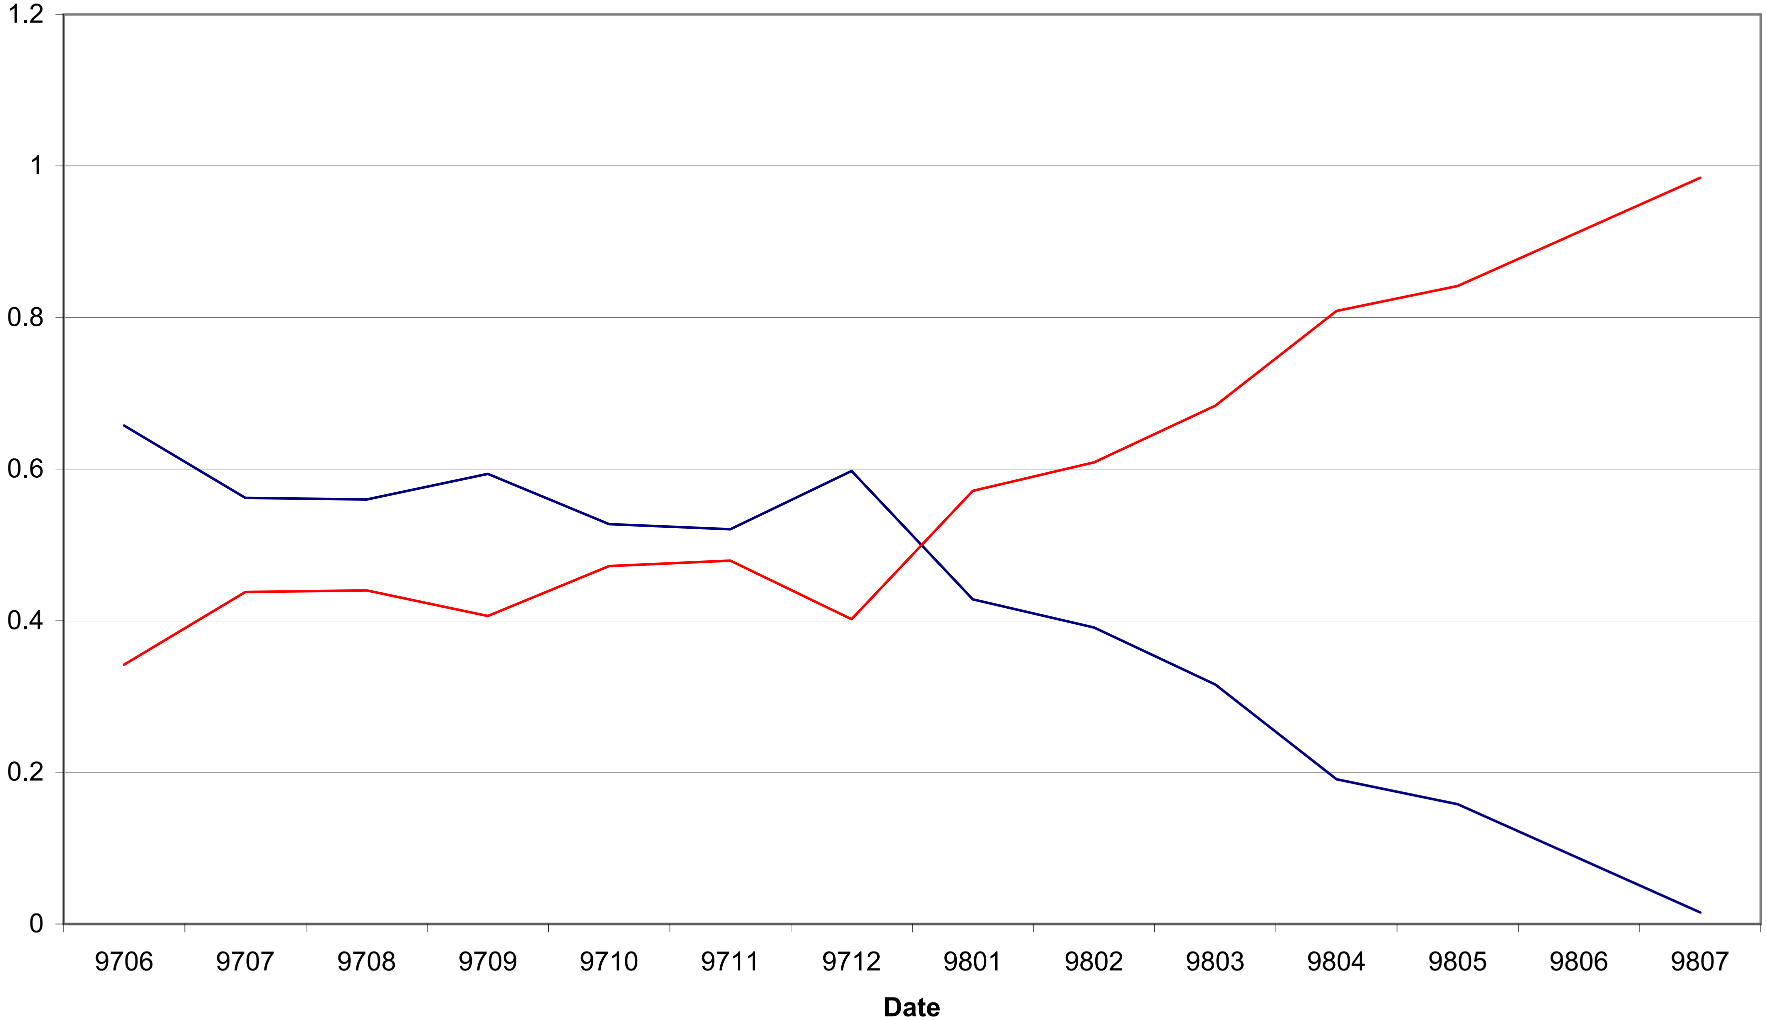

Figure 2
LIFFE & Eurex Price Impact Coefficients

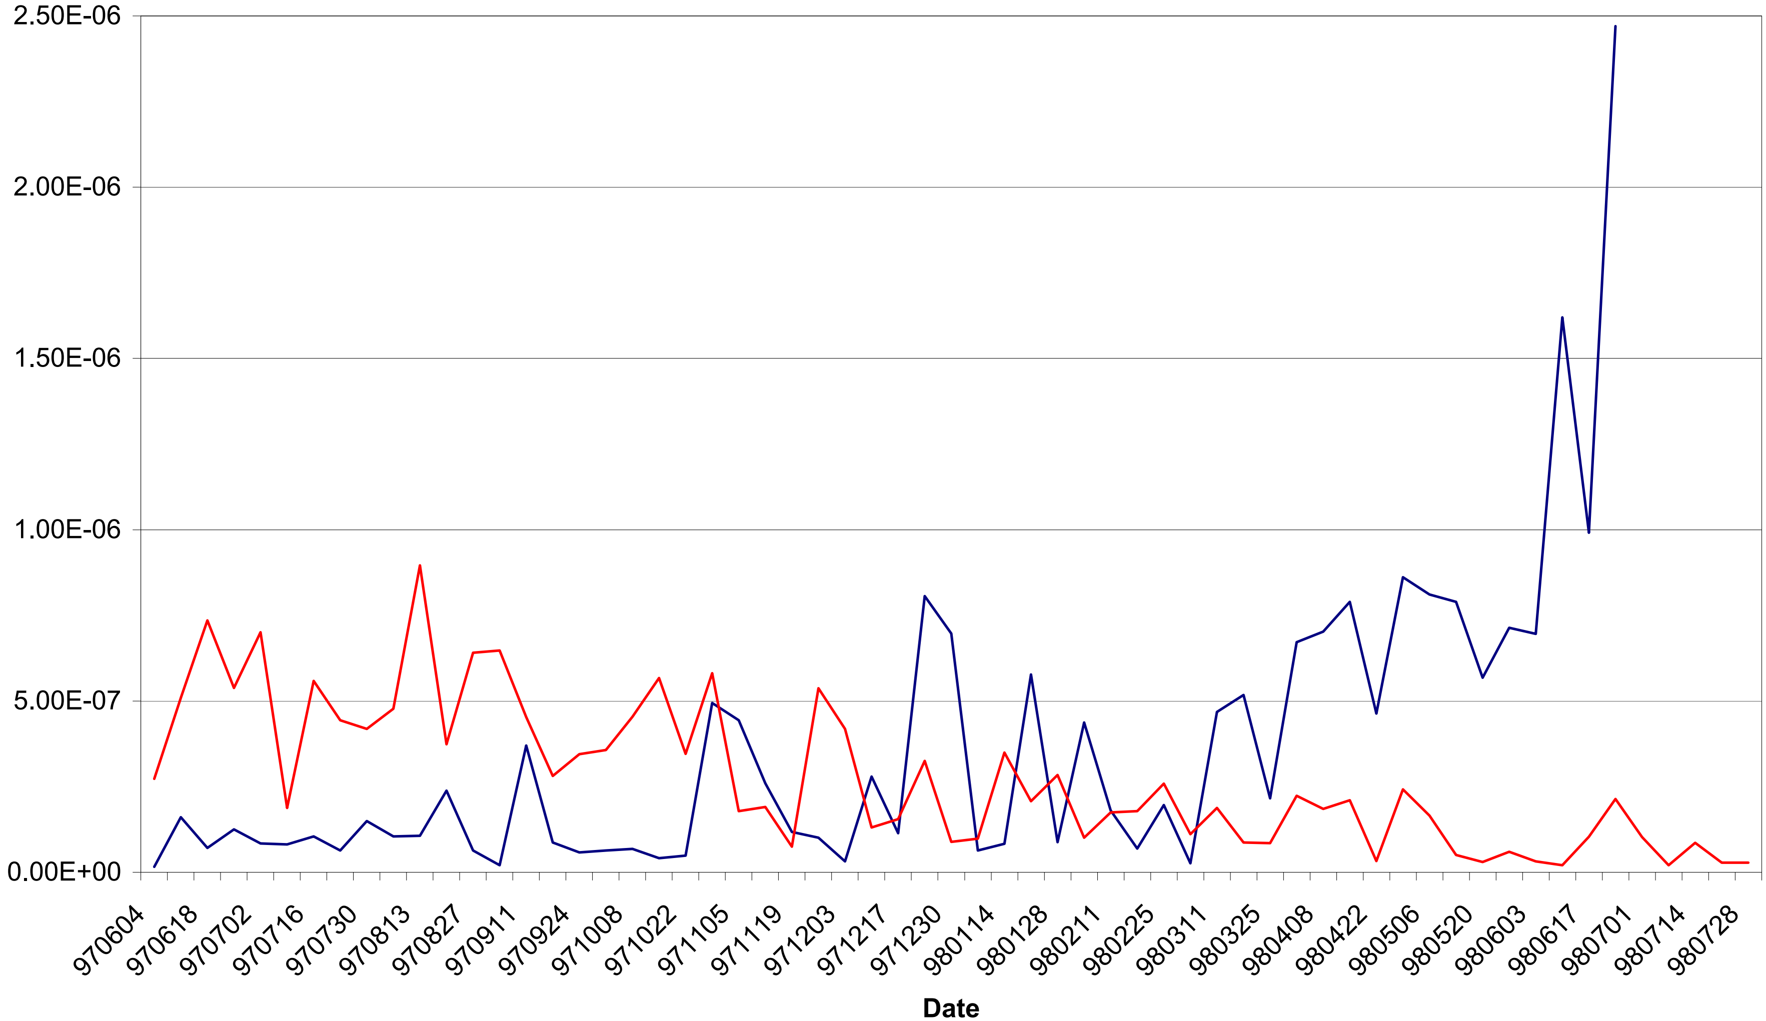

Figure 3
Price Impact Coefficient Volume
Combined LIFFE and Eurex Data