THE ROLE OF TRANSACTION COSTS FOR FINANCIAL VOLATILITY: EVIDENCE FROM THE PARIS BOURSE

Harald Hau
INSEAD

Abstract
This paper analyzes the causal linkage between transaction costs and financial volatility under two methodological improvements over the existing literature. First, we use panel data in which exogenous transaction cost differences in the French stock market are induced by price level dependent minimum price variation rules (tick size rules). Unlike in previous studies based on one-time regulatory tick size changes (like the U.S. decimalization), we can separately identify and control for marketwide volatility changes. Second, we avoid the pitfalls of biased volatility measurement across regimes by using the range as a tick size robust volatility metric. Panel regressions controlling for marketwide volatility effects show at high levels of statistical significance that the hourly range volatility of individual stocks increases by more than 30% for a 20% exogenous increase in transaction costs due to tick size variations in the French trading system. In the light of this evidence, higher transaction costs in general, and security transaction taxes in particular, should be considered as volatility increasing. (JEL: F3, G1, G14)

1. Introduction

> The introduction of a substantial government transfer tax on all transactions might prove the most serviceable reform available, with a view to mitigating the predominance of speculation over enterprises in the United States.
>
> —John Maynard Keynes, *The General Theory of Employment, Interest and Money*, 1936.

> Despite the prevailing opinion to the contrary, I am very dubious that in fact speculation in foreign exchange would be destabilizing. Evidence from some earlier experiences and from current free markets in currency in Switzerland, Tangiers, and elsewhere seem to me to suggest that, in general, speculation is stabilizing rather than the reverse, though the evidence has not

Acknowledgments: I thank seminar participants at INSEAD, Said Business School (Oxford), the Norwegian School of Management, the University of Oslo, and the University of Belfast for helpful comments. This research was supported by the Fondation Banque de France and Inquire Europe.
E-mail address: harald.hau@insead.edu

Journal of the European Economic Association June 2006 4(4):862–890

yet been analyzed in sufficient detail to establish this conclusion with any confidence.

—Milton Friedman, *Essays in Positive Economics*, 1953.

This paper provides new and robust evidence on a long and continuing debate about the relationship between trading costs and financial market volatility. At least since Keynes' stock market critique in 1936, stock price volatility has been related to low transaction costs which allegedly facilitate destabilizing financial speculation. Although the existing empirical evidence mostly suggests that higher transaction costs foster rather than mitigate financial price volatility, it suffers from serious methodological problems related to the data structure and biased volatility measurement. A new data structure combined with a refined volatility metric allows us to overcome these shortcomings and reach more robust and conclusive evidence on the issue.

The question about the nexus between transaction costs and financial volatility is interesting in at least three respects. First, regulatory, organizational and technological progress has considerably decreased transaction costs. Financial market liberalization in the 1980s lowered trading commissions and electronic trading in the 1990s further diminished stock trading costs.[1] At the same time, individual stock volatility appears to have increased in the U.S. (Campbell et al. 2001). It is unclear whether there is a causal link here or just coincidence. Second, transaction costs are influenced by the microstructure organization of the market. The introduction of smaller pricing grids (ticks) in the U.S. with price steps of 1/16th of a dollar instead of 1/8th appears to have reduced transaction costs for the majority of investors. The introduction of decimal quotation in 2001 further reduced transaction costs for small trades in the NYSE and Nasdaq (Bessembinder 2003). Does this regulatory transaction cost benefit come at the expense of higher stock price volatility, or do we obtain more price stability at the same time? Third, transaction costs sometimes include a tax component. Although security transaction taxes generally decreased in the 1990s, they remain, nevertheless, important in a few countries like the U.K.[2] Moreover, parts of the antiglobalization movement have elevated global security transaction taxes to one of their policy objectives. The policy debate about financial market stability seems to evolve around convictions rather than sound evidence.

Previous research on the nexus between transaction costs and stock price volatility suffers from two major shortcomings. First, many studies focus on one-time marketwide regulatory modifications of the tick size regime—like the introduction of decimal quotation. The analysis therefore lacks a proper panel data structure which identifies the policy event separately from a fixed time effect.

1. See for example Domowitz, Glen, and Madhavan (2001) and Jones (2002).

2. Stamp duties in the U.K. amount to an astonishingly high 0.5% of the transaction volume.

This data structure makes any volatility inference problematic under time varying volatility. A single intertemporal change in volatility can in principle explain the evidence. By contrast, our data features a panel structure in which intertemporal changes in volatility are separately identified and controlled for. The second and more important shortcoming of the existing literature concerns the volatility measurement itself. Volatility is measured by the standard deviation of midpoint returns, and this volatility metric is biased across tick size regimes.[3] Moreover, this volatility measurement bias goes in the direction of the "evidence." A larger pricing grid increases the dispersion of the rounding error between the true latent price and closest grid price. A larger dispersion of this rounding error implies a larger dispersion of the quoted midprice return and consequently a larger standard deviation measurement under large ticks. Previous findings in support of a lower standard deviation for the midprice return after a tick size decrease may represent a measurement artefact which is uninformative about the true underlying volatility change. By contrast, our own volatility inference is based on the range of the midprice. The average range constitutes a tick size robust volatility metric because return increasing and decreasing rounding errors cancel. Hence, a higher dispersion of the rounding error under larger ticks is without consequence for the average range measure. Unlike previous work, our paper provides both strong and robust evidence on the causal link between transaction costs and financial volatility.[4]

We use a large data set on the French stock transactions between 1995 and 1999 to show that higher transaction costs increase stock return volatility. During this period, French stocks were subject to an important transaction cost increase whenever their price moved above the French francs (FF) 500 price threshold. Above FF 500, the minimal tick size for quotes in the centralized electronic order book increased by a factor of 10 from FF 0.1 to FF 1. The smallest feasible percentage spread for stock quotation, therefore, increased from 2 to 20 basis points. We document that the 20 basis point spread is indeed frequently binding for stock prices above FF 500, and therefore constitutes an exogenous cost component induced by the pricing grid of the electronic order book. The inflated spread in the limit order book represents a cost analogous to a security transaction tax to speculators demanding liquidity. The large tick regime therefore renders speculation more expensive.

The data structure in our paper is different from those in previous event studies on regulatory modifications of the tick size regime. Marketwide tick size

3. See for example Bessembinder (2000, 2003) as well as Ronen and Weaver (2001). These authors use the standard deviation or variance as their volatility measure across different tick size regimes.

4. Compare also Alizadeh, Brandt, and Diebold (2002), who highlight the intriguing robustness of the range to microstructure noise.

modifications do not allow us to distinguish the volatility effect of a transaction costs change from a fixed time effect. However, in the French data, the tick size discontinuity occurs with respect to the stock price. It does not feature an intertemporal change in tick size regulation. The transaction cost effect is therefore separately identified relative to fixed time effects. French stock market regulation thus provides an ideal natural experiment on the role of transaction costs for stock return volatility.

Our sample selection consists of all CAC40 index stocks which trade in the price interval from FF 400 to FF 600 over the four-year period from January 1995 to December 1998. Effective spread measurements on approximately 4.7 million trades show that the median effective spread is 20% higher for stocks with prices just above FF 500. This finding is not new and corresponds to qualitatively similar results in the existing literature. For the same stock sample, we record the quoted midprice defined as the arithmetic average of the best bid and ask price. This midprice is used to calculate the percentage range defined as the difference between the highest and the lowest midprice relative to the midpoint between these two values. Unlike the standard deviation of returns used in previous work, the range provides a volatility measure which is unbiased across tick size regimes. Panel regressions show that the percentage hourly range is more than 30% higher for the stocks trading at prices above FF 500 after controlling for marketwide volatility effects. Our volatility inference is based on 47,213 hourly range measurements and the result is obtained at a high level of statistical significance.

Our evidence directly bears on the historical debate about the (de-)stabilizing role of short-term speculation. Higher transaction costs fall disproportionately on short-term speculators. Their speculative activity is clearly discouraged by higher transaction costs. But whether reduced speculation by short-term traders increases or reduces price volatility has always been controversial. Our evidence that lower transaction costs stabilize prices can therefore be interpreted as a rehabilitation of the short-term speculator. Reduced transaction costs increase his incentive for intertemporal speculation. Short-term speculation appears generally to be price stabilizing as conjectured by Friedman (1953), Miller (1991), and others.

The following section discusses the existing literature on the nexus between transaction costs and price volatility and the role of tick size regulation. We also explain how our results relate to the volatility effect of a security transaction tax. Section 3 introduces the institutional framework of the French stock market. We discuss in particular its tick size regime and the electronic trading system. Section 4 discusses the publicly available microdata and our sample selection. Methodological issues of spread and volatility measurement are discussed in section 5. Section 6 presents the empirical results for effective spreads and price volatility. Section 7 concludes.

2. Literature

2.1. Nexus between Transaction Costs and Volatility

The theoretical literature provides little guidance as to the relationship between transaction costs and financial price volatility. Some economists, such as Tobin (1978, 1984), Stiglitz (1989), Summers and Summers (1989), and Eichengreen, Tobin, and Wyploz (1995) conjectured that higher transaction costs discourage destabilizing investors with short-run horizons while being less costly for stabilizing investors with long-run horizons. Higher trading costs may privilege trading based on economic fundamentals. The opposing view is articulated by Friedman (1953), who argues that speculative behavior is generally price stabilizing irrespective of the time horizon. Miller (1991), Schwert and Seguin (1993), and Dooley (1996), among others, suggest that short-term speculation may be as beneficial as investment behavior based on a longer time horizon. The relative merits of these opposing views need to be judged in the light of the empirical evidence.

The evidence can be grouped into time-series studies, studies based on modified tick size regulation and panel studies. Early empirical work focuses on intertemporal transaction cost variations. Mulherin (1990) examines a long-run series of estimated trading costs in the NYSE and relates it to the daily volatility of the Dow Jones returns over the period 1897 to 1987. The data suggest a negative but statistically insignificant correlation. However, such long-run evidence is problematic because of parallel changes in the underlying market structure and possible measurement errors for the estimated transaction costs. Umlauf (1993) contributes an observation from the Swedish transaction tax experience in the 1980s. He finds that neither the introduction of a 1% round-trip transaction tax in 1984 nor its increase to 2% in 1986 decreased volatility in the Swedish stock market. However, the Swedish tax was collected from domestic security brokers and was increasingly avoided as a large percentage of trading volume in Swedish securities moved to international markets (Campbell and Froot 1994). Jones and Seguin (1997) report on the liberalization of mandated minimal commission rates in the U.S. This regulatory change decreased transaction costs in the NYSE and the AMEX markets in 1975. The authors find a reduction in the market volatility in the year following the deregulation, but the same volatility decrease, although less pronounced, was also registered for the previously unregulated Nasdaq market. Overall, the time-series evidence is at best weak.

A more powerful statistical strategy consists in the analysis of regulatory changes which concern the tick size regime of a stock markets. Ronen and Weaver (2001) find that the marketwide adoption of $1/16 ticks decrease return volatility (measured by the standard deviation) along with transaction costs. Similarly, Bessembinder (2003) documents that the introduction of decimal quotation in the NYSE and the Nasdaq in 2001 reduce both transaction costs and the standard deviation of midprice returns. Both studies claim therefore a positive linkage between transaction costs and financial price volatility. However, the volatility decreases after the tick size reduction can in principle be explained by an independent marketwide volatility decrease. Inference based on a single regulatory event remains therefore problematic.

Studies based on panel data should in principle provide the clearest evidence. An example is Bessembinder (2000), who examines the tick size discontinuity in the Nasdaq at $10 per share in 1995.[5] Stocks below $10 per share exhibit smaller percentage ticks and lower transaction costs. The data structure here is similar to our own data. Unfortunately, volatility is again measured by the standard deviation of midprice returns, thus rending the inference biased towards higher volatility under larger ticks. Finally, Bessembinder and Rath (2002) analyze stocks moving from the Nasdaq market to the NYSE. They find strong evidence that the newly NYSE listed stocks reduce both trading costs and the standard deviation of daily returns. But NYSE listings may simultaneously alter other volatility parameters related to a different market structure or investor composition. The cross-market comparison is therefore inconclusive because the volatility change may result from a stock listing effect and not from a transaction cost effect.

But the most important objection to the above evidence concerns the measurement methodology for volatility itself. All cited studies based on tick size effects measure volatility as the standard deviation of the midprice return. This volatility metric does not allow for robust volatility comparisons across different ticks size regimes. Intuitively, a larger pricing grid generates bigger rounding errors between the latent fundamental midprice the closest feasible midprice on the pricing grid. A larger dispersion of the rounding error implies a larger dispersion of the midprice return and larger standard deviation of the return for larger ticks. The evidence in the previous literature may therefore represent a measurement artefact due to a tick size sensitive volatility metric and be uninformative about the actual volatility change of the latent midprice. Panel data inference based on a tick size robust volatility metric is therefore needed to clearly establish a positive nexus between transaction costs and financial price volatility.

2.2. Ticks in the Literature

The literature on market microstructure has produced a large number of studies on the role of tick size for transaction costs. Two effects can be distinguished. First, bid-ask spreads may often come relatively close to the average tick size

5. The tick size discontinuity apparently comes from a quoting convention among market makers (and not an explict rule) to use 1/8th at or above $10 per share and 1/32th for bid quotations below $10 per share.

(Angel 1997). The minimal tick size regulation is therefore a frequently binding constraint and imposes exogenous differences in transaction costs. Second, higher tick sizes may decrease broker incentives for competitive quote improvement. Electronic trading systems generally reward quote improvements with privileged execution (price priority). The costs of such quote improvements in terms of price sacrifice is increased for a higher tick size. Harris (1994) predicts that larger tick size, therefore, reduces the incentive for competitive quote improvement and increases quoted spreads. Simultaneously, lower quoted spreads may reduce the incentive for liquidity provision. A smaller tick size may therefore result in lower market depth as measured by the liquidity offered at the best limit prices.

A large number of studies have confirmed the positive relationship between quoted spreads and percentage tick size for many different markets.[6] But transaction costs for large orders are influenced by both quoted spreads and market depth. A more meaningful measure of effective transaction cost is the volume weighted average of the execution prices (along the price elastic liquidity supply function) relative to the midprice. This so-called "effective spread" accounts for the price impact of larger market orders. Generally, results for effective spreads are qualitatively similar to those for quoted spreads. Bacidore (1997), Porter and Weaver (1997), and Ahn, Cao, and Choe (1998) confirm that the smaller tick size in the TSE reduced effective spreads by approximately 20%. For the NYSE move to 1/16th, Bollen and Whaley (1998) estimate an effective spread reduction of nearly 8%. Bessembinder (2000) studies Nasdaq stocks undergoing a tick size modification at the $10 price threshold and finds an 11% effective spread decrease due to smaller ticks.[7] Decimalization of the price quotes in the NYSE on January 29, 2001, and in the Nasdaq on April 9, 2001, reduced both tick size and effective spreads (Bessembinder 2003). Although transaction costs generally decrease along with the tick size, this benefit may mostly accrue to investors and speculators with small and medium size orders. For institutional investors with very large orders, the reduction in market depth may reduce or even outweigh the benefit of narrower spreads (Jones and Lipson 1999, 2000). But very large transaction volumes of institutional traders may not be the right benchmark for the short-term speculator. Given that transaction costs increase in volume, the marginal short-term speculator is likely to trade modest quantities, since he is free to choose the size of his position. He should therefore unambiguously reduce his trading costs in a trading environment with smaller ticks.

6. See for example Lau and McInish (1995); Ahn, Cao, and Choe (1996, 1998); Ronen and Weaver (2001); Bacidore (1997); Porter and Weaver (1997); Bollen and Whaley (1998); Goldstein and Kavajecz (2000); Jones and Lipson (2000).

7. Unlike in the Paris Bourse, the Nasdaq tick size change from 1/32th to 1/8th at prices of $10 is based on a market convention rather than imposed by the trading system.

2.3. Tick Size Effects and Security Transaction Taxes

Can tick size effects serve as an experiment to evaluate the volatility effect of a security transaction tax? On the liquidity demand side, it certainly makes no difference if the transaction cost increase originates in tick size regulation or in a security transaction tax with the same spread increase. Hence, demand side effects are equivalent. However, the same does not hold for the liquidity suppliers or brokers, for whom a security transaction tax is different from a binding tick size constraint. While a tax is a rent for the tax authority, binding tick size regulation constitutes a rent for the liquidity suppliers.

The latter makes liquidity provision more profitable and may generate a more liquid market. This results in higher market depth documented by Goldstein and Kavajecz (2000) for the NYSE and Ahn, Cao, and Choe (1998) for the Toronto Stock Exchange.[8] Greater market depth should generally reduce volatility because of a lower price impact of large market orders. The positive liquidity supply effect of a tick size increase is absent if the larger spread is induced by a security transaction tax. In this case, the liquidity provision (through limit order submission) itself is subject to taxation, and no increase in liquidity provision can be expected. These considerations lead us to conclude that security transaction taxes generate more price volatility than binding tick size regulation for a similar increase in spreads. The volatility effects of higher transaction costs estimated in our study should therefore be interpreted as a lower limit for the volatility increase due to a security transaction tax. Finally, we highlight that a comprehensive discussion on the Tobin tax includes many aspects outside the scope of this paper. For a general discussion on the Tobin tax we refer to ul Haq, and Grunberg Kaul, (1996). But it is fair to say that the linkage between transaction costs and volatility is at the core of the theoretical debate.

3. Institutional Framework

Since the beginning of the 1990s, the Paris Bourse has operated as a computerized and centralized limit order market. It allows for continuous trading from 10:00 a.m. to 5:00 p.m.[9] The opening price at 10:00 a.m. is determined by a preopening mechanism with an initial auction (see Biais, Hillion, and Spatt 1999). All brokers with trading terminals enjoy equal trading opportunities in the computerized system known as CAC (Cotation Assistée en Continue). There are no market makers or floor traders with special obligations.

8. By contrast, Ronen and Weaver (2001) find no change in market depth related to the adoption of 1/16th in the AMEX in May 1997.

9. A final batch auction after 5:00 p.m. was introduced on June 2, 1998. We also note that since April 21, 2002, the trading period ranges from 9:00 a.m. to 5:30 p.m.

3.1. The Tick Size Regimes

Investors can submit limit orders at any price on a prespecified pricing grid, defined by the tick size. This tick size, that is, the minimum price step between two prices accepted by the trading system, depends upon the price level of the security. For prices below FF 5, the tick size is FF 0.01; for prices between FF 5 and FF 100 the tick size is FF 0.05; for prices between FF 100 and FF 500 the tick size is FF 0.1; for prices between FF 500 and FF 5,000 the tick size is FF 1; and above FF 5,000 the tick size is FF 10. During the sample period 1995–1998, the French franc was worth approximately \$0.18. Most stocks are traded in the price range from FF 200 to FF 1,000 and are therefore subject to either FF 0.1 ticks (referred to as small ticks) or FF 1 ticks (referred to as large ticks). The value of FF 500 marks an empirically important discontinuity in the electronic order book at which the price grid increases by a factor of 10.[10] The pricing grid imposes a technical lower bound on the smallest possible percentage spread between the best bid and ask price. Spreads cannot decrease below 20 basis points for stocks just above FF 500, although they can drop to 2 basis points below a security price level of FF 500. We show in section 5 that the minimum spread imposed by the large tick regime is indeed frequently binding and therefore artificially inflates investor trading costs.

Statistical inference based on the step function of the pricing grid provides a better natural experiment compared to marketwide tick size reform, which subjects all stock (or entire stock groups) to a one-time tick size modification. By contrast, grid size step functions imply that the spread constraint operates on a random subsample of stocks with the unconstrained stocks available as a control group. We can therefore distinguish the transaction cost effect on volatility from other market wide-volatility shocks.

The tick size regime of the French stock market was modified with the introduction of euro quotations on January 2, 1999.[11] The minimum percentage spread can no longer exceed 10 basis points compared to 20 basis points before 1999. This suggests that we focus our empirical analysis on the period prior to the introduction of the euro when tick size regulation was more likely to impose a constraint on the quoted spread.[12]

10. By comparison, the tick size jump from 1/32th to 1/8th at \$10 in the Nasdaq market prior to 1997 is based on an informal market convention rather a rule imposed by the trading system. Moreover, it concerns mostly small and illiquid stocks. Also, the NYSE tick size breakpoint at 1 dollar is irrelevant for most stocks.

11. The new euro price grid was designed to limit the maximal percentage ticks size at 10 basis points. Tick size for prices below €50 is €0.01; for prices between €50 and €100 the tick size becomes €0.05; for prices between €100 and €500 the tick size is €0.1; and for prices above €500 the tick size is €0.5. The empirically most relevant tick size discontinuity at around €50 is now reduced to a grid factor of 5.

12. Compare Bourghelle and Declerck (2004) for a study on market liquidity changes at the Paris Bourse related to the transition to the euro tick regime.

3.2. The Trading System

Like most electronic markets, the Paris Bourse enforces price and time priority. Orders are executed at the best available price. If two limit orders offer the same price, execution preference is given to the limit order which arrives first. The electronic order book itself is very transparent. Information on the five best bid and ask prices and the number of shares demanded or offered at each of these prices is continuously available to the public. Brokers can observe the entire limit order book and the identification codes of the brokers placing orders.[13]

One specificity of the trading system is the treatment of "market orders" without a limit price. The CAC trading system treats them automatically as limit orders at the best momentarily available price. Execution is therefore partial if the demand exceeds the available liquidity at the best price. The nonexecuted fraction of such an order is transformed into a limit order. However, traders can always obtain full execution by selecting a sufficiently unfavorable limit price.

Essentially, all trades are executed at prices in the electronic book, except prematched block trades, which are subject to special rules. If the prematched block trades occur at or inside the current spread, they can bypass the limit order book. If they occur outside the current spread, the priority of the previously posted limit order is respected. For example, if the block price exceeds the best ask, the limit orders between the best ask and the block price are purchased by the block buyer at the block price. Approximately 1.1% of the trades and 17.2% of the volume occurs through prematched trades.

By law, the French stock market is a centralized market. Transactions governed by French contracts must be executed on the Bourse. Trading outside France is of course possible. Dealers in London may, for example, bypass the Paris market by using the London International Stock Exchange Automated Quotation System (SEAQ) to search for counterparties with a trading interest. De Jong, Nijman, and Röell (1995) document with a short data sample in 1991 that this happens particularly for large trades. The London transaction prices are negotiated between dealers and are not subject to formal tick size constraints. We assume that such interdealer trades outside France do not substantially modify the transaction cost pattern induced by the tick size regulation in the main Paris market. Robustness of this assumption can be checked by excluding those stocks for which more liquid parallel markets exist. We identify all sample stocks with cross listings in the London Stock Exchange or the New York Stock Exchange (ADRs) during the period 1995 to 1999. This is the case of 6 out of 26 stocks in

13. An exception to full order book transparency are so-called "hidden orders". These are orders for which only a fraction of the available liquidity appears on the trading screen. For an analysis of the role of hidden orders at the Paris Bourse, see Harris (1996).

the sample.[14] Parallel trading is presumably strongest in these stocks. However, their exclusion from the analysis did not quantitatively alter the results.

4. Data and Sample Selection

The Paris Bourse publicly provides comprehensive historical microdata on best limit quotes and security transactions.[15] Our data selection is motivated by two concerns. First, tick size regime induced transaction cost differences are likely to be most relevant for large and highly liquid stocks. These tend to have relatively small transaction costs and the tick size constraint for spread quotation is more frequently binding. We therefore limit our analysis to the stocks in the CAC40 index comprising the 40 largest and most liquid French stocks. CAC40 stocks account for approximately 64% of all transactions in our data period. Second, the transition to the euro quotation of stock prices in 1999 also brought a modification of the tick size regime toward a smaller tick size. Statistical identification of an exogenous transaction cost effect is therefore better assured by using data prior to January 1999. We focus our analysis on four years of microdata from January 1995 to December 1998.

Nevertheless, four years of quote and transaction data for all CAC40 stocks still exceeds our data processing possibilities. We therefore choose to observe only those CAC40 stocks which are quoted in a price window around the tick size discontinuity at FF 500, namely between FF 400 and FF 600. The tick size constraint for the minimal percentage spread is obviously most severe directly above FF 500 and least so directly below. As stock prices move away from the FF 500 threshold, the two tick regimes become more similar in terms of their minimal feasible percentage spread. For example, a tick size of FF 1 at a stock price of FF 1,000 allows for a 10 basis point percentage spread, just as a tick size of FF 0.1 for a stock price of FF 100. The choice of a relatively small price window from FF 400 to FF 600 limits the number of observations and focuses on those observations for which the tick size regulation is most discriminatory.

All data are obtained directly from the Paris Bourse on monthly CD-ROMs which combine a variety of data files on transactions and quotes in different market segments. We match two of these files to calculate effective spreads for individual trades. A first data file (coded BDM2D2) provides a continuous record of the best bid and ask price of every stock. These data allow us to construct a continuous midprice as the benchmark for the transaction prices. A second data

14. The London cross listings occured for Lafarge on October 30, 1972, for Total on September 26, 1973, for Saint-Gobain on July 2, 1987, and for Alcatel Alstom on June 25, 1998. NYSE trading in ADRs prior to 1999 was feasible for Alcatel Alstom, AXA, France Telecom, and Total.

15. Previous studies on the same data source include Biais, Hillion, and Spatt (1995, 1999), and Venkataraman (2001).

file (coded BDM1D2) contains a complete record of all trades and subtrades stripped of the identity of the counterparties. A market order executed against various limit orders is documented with the corresponding number of subtrades. By matching the transaction price with the midprice at the transaction time, we calculate the effective percentage spread for each trade and the trade weighted effective spreads for each executed order. We also use a third data file (coded BDM5D2) with records of the index level for the CAC40 index every 30 seconds. This allows us to calculate index volatility.

Unfortunately, the data on the best quotes do not contain any information about the best bid and ask price during the opening auction. Registration of the best quotes only starts with the first transaction in the regular continuous trading period. We can therefore only calculate a midprice shortly after the 10:00 a.m. opening auction. Spread calculations on transactions in the opening auction are therefore difficult. These transactions are ignored in the consecutive analysis.[16]

Because our data selection criterion is based on both a stock belonging to the CAC40 index and a particular price range, it is useful to first provide an overview of the resulting stock sample. Table 1 presents summary statistics separately for stocks in the price range from FF 400 to FF 500 (small ticks) and from FF 500 to FF 600 (large ticks). A total of 26 stocks in the index trade (in terms of their hourly range midprice) between FF 400 to FF 600 for at least 20 trading hour between January 1995 and December 1998. Altogether, we obtain 47,213 trading hours for 26 stocks. Two stocks trade for less than 20 hours in the respective price interval, namely Schneider and Cap Gemini. We exclude their negligible number of volatility observations to obtain a more balanced data panel. Most sample stocks (21 out of 26) trade in both tick size regimes, while 1 stock trades exclusively in the small tick regime, and 4 trade only in the large tick regime. On average, stocks are recorded in the small tick regime for 1,067 hours and in the large tick regime for 949 hours. Overall, we obtain 23,481 stock trading hours in the small and 23,732 stock trading hours in the large tick regime. The average number of daily trades are 779 and 652 for the small and large ticks, respectively. Average daily volumes are FF 107 million and FF 110 million, respectively.

5. Methodology

This section discusses the transaction cost and volatility measurement. We focus in particular on the measurement bias related to tick size sensitive volatility measures like the standard deviation. This leads us to advocate a robust volatility metric given by the range. Our statistical inference then consists in a straightforward

16. We also filter the data for outliers. Transactions for which the quoted spread exceeds 10% or is negative are discarded.

TABLE 1. Stock characteristics by tick regime.

	Small ticks				Large ticks			
		Daily averages of				Daily averages of		
Stock name	Trading hours	Trades	Volume	Market cap.	Trading hours	Trades	Volume	Market cap.
Sanofi	609	514	81	46.6	1,460	651	108	58.1
Total	1,033	767	168	106.8	582	1,444	406	138.8
Elf Aquitaine	646	2,444	234	116.7	630	2,044	330	151.4
Bouygues	753	345	32	11.5	3,785	293	33	12.6
Lyonnaise des Eaux	2,886	366	46	26.8	1,441	576	97	32.3
Lafarge	945	914	104	41.6	627	909	123	54.7
AXA	615	1,699	360	142.7	269	1,971	467	189.8
CFF	52	304	13	5.9	336	221	11	6.7
BIC	1,230	539	45	25.5	96	480	37	24.1
Bancaire (CIE)	629	256	24	12.5	2,661	260	27	14.7
Général des Eaux	1,091	562	85	55.0	2,016	609	120	63.2
Spie Batignolles	404	1,088	161	70.5	58	831	120	78.0
Paribas	1,165	1,247	188	59.5	461	1,628	342	84.4
CCF	869	624	85	32.7	288	573	94	38.0
Alcatel Alstom	3,201	957	144	68.3	486	2,905	486	87.9
Havas	2,291	445	59	28.0	275	631	225	44.3
Valéo	741	613	73	33.7	457	506	70	39.0
Société Général	110	695	88	40.9	3,255	606	97	49.0
Crédit LCL France	2,758	345	32	15.8	1,256	493	57	20.6
BNP	396	2,021	292	96.9	363	1,819	313	112.1
SGS Thomson	648	725	117	59.8	115	654	106	72.4
France Télécom	409	4,165	503	425.3	—	—	—	—
Accor	—	—	—	—	799	262	38	14.5
Peugeot	—	—	—	—	1,181	443	71	28.9
Saint-Gobain	—	—	—	—	711	590	90	47.7
Canal Plus	—	—	—	—	124	244	15	12.7
All Stocks	23,481	779	107	53.4	23,732	652	110	44.6

Notes: Reported are trading hours; the average number of daily trades (and subtrades), average daily volume (in millions of French francs, FF), and average market capitalization (in billions of FF) for all stocks in the French CAC40 index which trade with prices between FF 400 and FF 600 for at least 20 hours over the four-year period from January 1995 to December 1998. We distinguish hourly and daily trading periods according to the tick size regime to which each stock is subject. Stocks quoted in the price range from FF 400 to FF 500 are subject to a minimum tick size of FF 0.1 (small ticks), and stocks in the price range from FF 500 to FF 600 are subject to a minimum tick size of FF 1 (large ticks).

comparison of transaction costs and range volatility for stock observations in the two tick size regimes. Finally, we refine the volatility analysis with panel regressions controlling for stock specific effects, intraday volatility patterns, and volatility autocorrelation.

5.1. Transaction Cost Measurement

The transaction cost measurement follows standard conventions. For individual trades and subtrades, we calculate the effective spreads as twice the distance from the midprice measured in basis points. For a transaction price P^T and a midprice P^M as the arithmetic average of the best bid and ask price, we obtain the effective spread (in percentage terms) as

$$\text{SPREAD}^{\mathit{Trade}} = 200 \times \frac{|P^T - P^M|}{P^M}.$$

To obtain a more Gaussian distribution, we use a (natural) logarithmic transformation and define the log effective spread as

$$\text{LSPREAD}^{\mathit{Trade}} = \ln\left[\text{SPREAD}^{\mathit{Trade}}\right].$$

Alternatively, we can measure transaction costs for executed orders. A single order might be executed in n subtrades against limit prices $P_1^T, P_2^T, \ldots, P_n^T$ with corresponding quantities $V_1, V_2, \ldots, V_n$. We denote the executed order volume as $V = \sum_{i=1}^{n} V_i$. The effective transaction price follows as the value weighted average of the traded prices, $\bar{P}^T = \sum_{i=1}^{n} P_i^T V_i / V$, and the effective spread of an executed order (in basis points) is defined as

$$\text{SPREAD}^{\mathit{Order}} = 200 \times \frac{|\bar{P}^T - P^M|}{P^M}.$$

The log effective spread again follows as

$$\text{LSRPREAD}^{\mathit{Order}} = \ln\left[\text{SPREAD}^{\mathit{Order}}\right].$$

We highlight the fact that the effective spread for orders measures the transaction costs only with respect to a single transaction. Brokers might break large client orders into many smaller orders for consecutive execution. These multiple transactions are likely to result in higher transaction costs than those measured by the effective spread because of a consecutive price impact. But the transaction data of the Paris Bourse do not allow us to identify transaction sequences pertaining to the same broker. The effective spreads are therefore the best available transaction cost measure.[17]

17. For a transaction cost analysis of large institutional traders, see Jones and Lipson (1999, 2000).

5.2. Pitfalls of Volatility Measurement based on the Standard Deviation

Accurate measurement of stock price volatility is crucial for our analysis. The French stock market provides us not only with a record of all transaction prices, but also with data on the best bid and ask price. Best bid and ask quotes can be used to calculate the midprice as the arithmetic average throughout the trading day. Price measurement based on the midprice of the best quoted bid and ask price alleviates so-called bid-ask bounce effects inherent in the transaction prices. Such bid-ask bounce effects are likely to depend on the minimum tick size and would therefore distort the volatility comparison across tick size regimes. But even a volatility measure based on the midprice may not be robust to differences in the tick size. This is important since we compare volatility measurements across different tick size regimes. For example, the standard deviation of midprice returns suffers from a tick size distortions and therefore does not constitute a suitable volatility measure in our experiment.[18]

To illustrate this point, we define P_t the true latent price and $\Xi(P_t)$ quotable price rounded to the closed available full tick. Let $e_t = (\Xi(P_t) - P_t)/P_t$ denote the percentage rounding error. For a percentage tick size *PTS*, the maximal percentage rounding error is 1/2(PTS). We assume that the initial unconditional price distribution is uniform. The percentage rounding error e_t is then also uniformly distributed over the interval $[-1/2(PTS), 1/2(PTS)]$. Furthermore, assume a latent price process with zero expected return and a variance $\mathrm{Var}(R_t)$ over the measurement interval. The variance of measured returns $\widetilde{R}_t = [\Xi(P_{t+1}) - \Xi(P_t)]/\Xi(P_{t+1})$ can be approximated as

$$\mathrm{Var}(\widetilde{R}_t) \approx \mathrm{Var}(R_t) + E[e_{t+1}]^2 - 2E[e_{t+1}e_t] + E[e_t]^2.$$

It is straightforward to determine the measurement bias of the variance for the case of large price movement for which the two consecutive measurement errors are approximately uncorrelated, hence $E[e_{t+1}e_t] \approx 0$. The variance of a uniformly distributed rounding error is given by $E[e_{t+1}]^2 = E[e_t]^2 = 1/12(PTS)^2$. The measurement bias of the standard deviation follows as

$$\begin{aligned}\text{Measurement Bias} &= \frac{SD(\widetilde{R}_t) - SD(R_t)}{SD(R_t)}\\ &= -1 + \sqrt{\frac{\mathrm{Var}(\widetilde{R}_t)}{\mathrm{Var}(R_t)}} \approx -1 + \sqrt{1 + \frac{\frac{1}{6}(PTS)^2}{\mathrm{Var}(R_t)}}.\end{aligned}$$

The measurement bias for the standard deviation is always positive and increasing in the percentage tick size *PTS*. It is also sensitive to the variance of the latent

18. The same tick size sensitivity applies to realized volatility defined as the sum of squared midprice returns. For a discussion of realized volatility see Andersen et al. (2001).

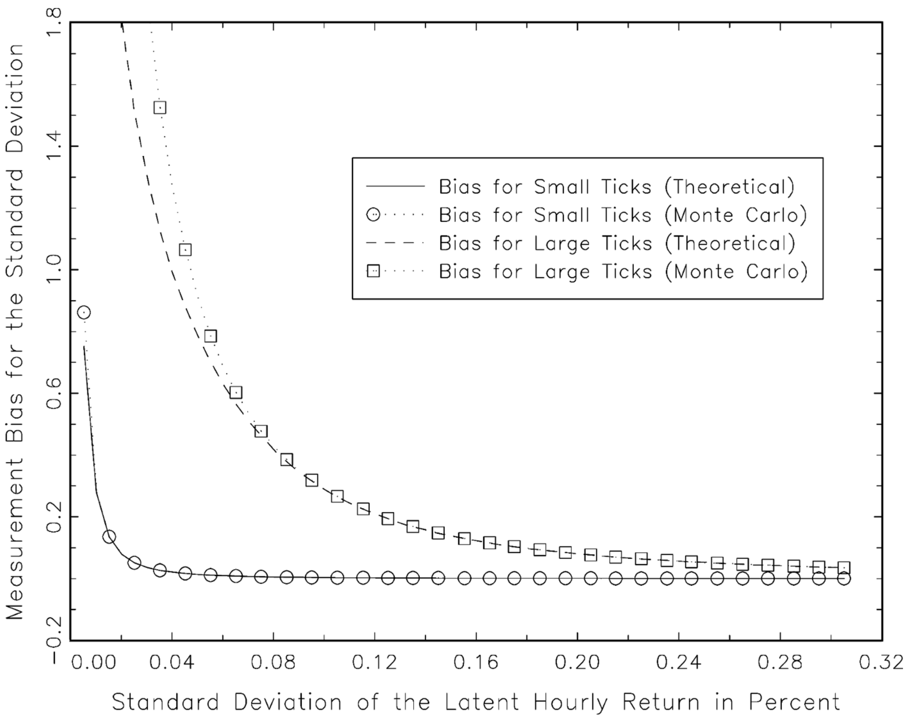

FIGURE 1. Measurement bias for the standard deviation as a function of the latent hourly return volatility under small and large ticks of 2 and 20 basis points, respectively.

return process over the measurement interval. For a low variance of the latent return or short measurement intervals, the bias can be become large for any given tick size. But the approximation $E[e_{t+1}e_t] \approx 0$ becomes less accurate if the return volatility decreases as the autocorrelation of the rounding error increases.

To evaluate the accuracy of the derived measurement bias, we use a Monte Carlo method with 1 million random draws from an initially uniform distribution and assume a normally distributed return innovation. The standard deviation of the true (latent) return is compared with measured return under the small and large tick size with $PTS = 0.1/500$ and $PTS = 1/500$, respectively. Figure 1 compares the theoretical measurement bias with the bias implied by the Monte Carlo simulations. The theoretical and simulated bias are very close for a standard deviation of the hourly return above 0.05%. Moreover, the difference in the measurement bias is of considerable economic magnitude in the large tick size regime for return volatility below a standard deviation of 3%. We note that the median hourly volatility in our French equity data is only 1.2%, which implies a volatility bias of approximately 23% for the large ticks and only 2% for the small ticks. Any volatility comparison across the two tick size regimes based on the standard deviation as the volatility metric would therefore be very misleading. It is also clear that the measurement bias is highly nonlinear in the underlying return volatility, which makes any posterior bias correction very difficult.

Generally, volatility measurement at high frequency should avoid a volatility metric which is convex in the rounding error. It is this convexity which induces an upward tick size dependent measurement bias as return increasing rounding errors enter the volatility metric more strongly than return decreasing rounding errors. The following section proposes a tick size robust volatility estimator.

5.3. Robust Volatility Measurement based on the Range Estimator

The so-called range represents a volatility metric which is not only tick size robust, but also highly efficient as shown by Alizadeh, Brandt, and Diebold (2002). The mean absolute return deviation is an alternative tick size robust volatility metric, but much less efficient compared to the range. The absolute return metric uses only the first and the last observation of a measurement interval, but ignores all observations within this interval. The range, on the other hand, scrutinizes all observations within the interval (for the smallest and largest value) and therefore uses more information than the absolute deviation. Moreover, the natural logarithm of the range features a near Gaussian distribution, and therefore represents a particularly attractive volatility measure for panel regressions. Formally, we define the percentage range (in percentage terms) as the difference between the highest and lowest midprice over a fixed interval I_t relative to their range midpoint, hence

$$\mathrm{RANGE}_t = 200 \times \frac{\max_{\{s \in I_t\}} P_s^M - \min_{\{s \in I_t\}} P_s^M}{\max_{\{s \in I_t\}} P_s^M + \min_{\{s \in I_t\}} P_s^M}.$$

The range measure can be zero for a quoted midprice which is constant over the respective time interval. To allow for a log transformation, we add a small constant to the range which corresponds to 10 basis points.[19] The log range is then defined as

$$\mathrm{LRANGE}_t = \ln\left[0.1 + \mathrm{RANGE}_t\right].$$

For our sample period, the continuous auction market at the Paris Bourse operates from 10:00 a.m. to 5:00 p.m. We choose every full trading hour as the sampling interval for the range measure. The hourly log range is calculated for all stocks in the CAC40 index which trade (with respect to their hourly range midpoint) in the price interval from FF 400 to FF 600 for at least 20 trading hours over

19. This is close to the median effective spread for large stocks in the small tick regime.

the 4-year period from January 1995 to December 1998. The data sample consists of 47,213 hourly range measurements for 26 different stocks.

6. Evidence

6.1. Transaction Cost Evidence

Transaction costs can be measured with respect to either trades or executed orders. In the latter case, we group all subtrades resulting from the same order into one single transaction. We count a total of 4,696,422 trades and 2,918,829 executed orders. Excluded in this count are pre-matched trades (1.15% of all trades) and all trades in the opening auction for which we cannot calculate the midprice (8.95% of all trades).

Table 2, panel (A), summarizes the distribution of the log effective spread by tick size regime for all trades. We provide a 1% confidence interval for the various centiles of both spread distributions using the binomial-based method. All reported centiles differ signigicantly for the two spread distributions. The median (mean) log effective spread in the small tick regime is -1.849 (-1.923) compared to -1.667 (-1.371) for large ticks. The median effective spread increase, therefore, amounts to 20% ($= e^{-1.667+1.849} - 1$) for trading with large ticks. Table 2, panel (B), reports the corresponding effective spread statistics for executed orders. As in the case of individual trades and subtrades, we find that the percentiles of the distribution are very different across the two tick regimes with a strong censoring effect on the left tail in the large tick regime. The latter shows a skewdness of 1.485 and 1.528 for trades and executed orders, respectively.

It is instructive to visualize the distribution of log effective spreads. Figure 2 plots the log effective spread for a random sample of 20,000 trades as a function of the price level. If effective spreads were plotted directly, their y-axis value would often coincide because of the discreetness of the ticks. To avoid this clustering effect, we add a small amount of random noise to each spread observation. This makes individual spread observations visually distinct, and the feasible spreads appear as a narrow band of points instead of a line. More points and a darker band show a higher density of spread observations. The lowest band in Figure 2 corresponds to effective spreads of 2 basis points ($\ln(0.02) = -3.91$) for tick steps of FF 0.1 below FF 500. The following band corresponds to a 4 basis point spread, and so forth. Continuity of the band over the entire price range from FF 400 to FF 600 is only reached with the 20 basis point spread band ($\ln(0.20) = -1.61$). Hence, we can clearly visualize that the tick size regulation is frequently binding for stock prices below FF 500. Figure 3 provides a nonparametric kernel density estimation of the effective spread of 2,540,764 trades below and 2,155,658 trades above FF 500. The two density distributions of the effective spread are indeed

TABLE 2. Log effective spreads for trades and executed orders by tick regime.

Panel A: Effective spreads for trades and subtrades

	Small ticks			Large ticks		
Percentiles	Centiles	1% C.I.		Centiles	1% C.I.	
1%	−3.8669	−3.8669	−3.8669	−1.7893	−1.7893	−1.7893
5%	−3.7567	−3.7567	−3.7564	−1.7758	−1.7758	−1.7758
10%	−3.1551	−3.1555	−3.1547	−1.7639	−1.7639	−1.7639
25%	−2.4753	−2.4761	−2.4744	−1.7272	−1.7272	−1.7272
50%	−1.8488	−1.8499	−1.8483	−1.6668	−1.6668	−1.6668
75%	−1.3005	−1.3009	−1.2997	−1.0219	−1.0225	−1.0207
90%	−0.7602	−0.7608	−0.7595	−0.5869	−0.5869	−0.5869
95%	−0.4340	−0.4350	−0.4340	−0.2983	−0.3001	−0.2983
99%	0.1522	0.1520	0.1534	0.2886	0.2886	0.2894
Observations		2,540,764			2,155,658	
Mean		−1.923			−1.371	
Std. dev.		0.944			0.528	
Skewdness		−0.157			1.465	
Kurtosis		2.791			4.702	

Panel B: Effective spreads for executed orders

	Small ticks			Large ticks		
Percentiles	Centiles	1% C.I.		Centiles	1% C.I.	
1%	−3.8713	−3.8713	−3.8713	−1.7893	−1.7893	−1.7893
5%	−3.7680	−3.7683	−3.7680	−1.7775	−1.7775	−1.7775
10%	−3.2040	−3.2050	−3.2032	−1.7639	−1.7639	−1.7639
25%	−2.6119	−2.6122	−2.6117	−1.7308	−1.7308	−1.7308
50%	−1.9815	−1.9826	−1.9803	−1.6743	−1.6743	−1.6743
75%	−1.4598	−1.4600	−1.4596	−1.0664	−1.0664	−1.0664
90%	−0.9331	−0.9337	−0.9325	−0.9243	−0.9254	−0.9243
95%	−0.6625	−0.6630	−0.6615	−0.5794	−0.5800	−0.5794
99%	−0.1338	−0.1338	−0.1332	−0.1017	−0.1017	−0.0998
Observations		1,587,683			1,331,146	
Mean		−2.046			−1.455	
Std. dev.		0.892			0.428	
Skewdness		−0.213			1.528	
Kurtosis		2.692			4.761	

Notes: The distribution of the log effective spread is provided separately for trades and subtrades (in Panel A) and entire market orders (in Panel B) for all stocks in the French CAC40 index with a price range (in French francs, FF) from FF 400 to FF 600 over the four-year period from January 1995 to December 1998. A market order can be partially excuted against various limit orders resulting in multiple subtrades. All effective spreads are calculated separately for stocks quoted in the price range from FF 400 to FF 500 subject to a minimum tick size of FF 0.1 (Small ticks), and stocks in the price range from FF 500 to FF 600 subject to a minimum tick size of FF 1 (Large ticks).

very distinct. Low spread density peaks occur below 20 basis points only for the small tick regime. For the large tick regime the density peaks with the first feasible spread, indicating the censoring effect of the tick constraint. Figure 4 provides the analogous density plot for the effective spread on executed orders, which closely resembles the corresponding plot for trades.

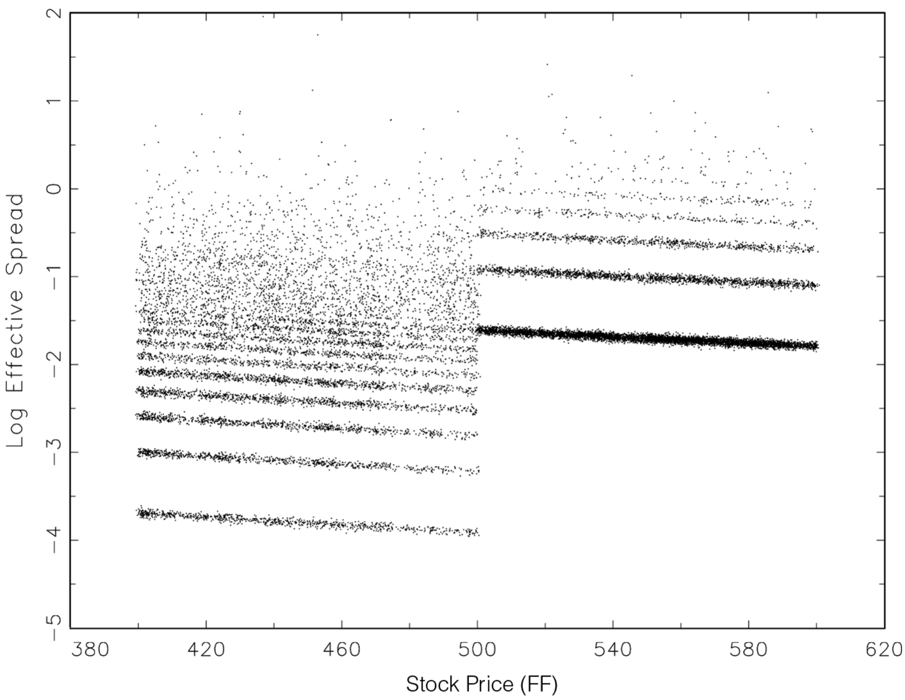

FIGURE 2. The log effective spread is plotted for a random sample of 20,000 trades on stocks in the price range from FF 400 to FF 600. At FF 500, the minimal tick size increases from FF 0.1 to FF 1. A small amount of noise is added to each observation to render it visually distinguishable.

These results clearly show that the tick size constraint in the French stock market is frequently binding for CAC40 index stocks with prices above FF 500, and comes with a statistical and economically significant transaction cost increase. Based on this exogenous transaction cost identification, we can now proceed to explore the volatility implications.

6.2. Volatility Evidence

The continuous price record of the Paris Bourse allows us a very precise volatility measurement. Table 3 provides the summary statistics on the hourly log range for each individual stock in the sample set by tick regime. Of the 21 stocks subject to both tick regimes, 19 show a higher mean for the log range in the large tick regime. Overall, the mean log range (LRANGE) is −0.542 for small ticks compared to −0.423 for large ticks. The skewness and kurtosis parameters indicate that the distribution of the log range is approximately normal for both the small and large tick subsamples. Table 4 states the volatility centiles by regime. The large tick

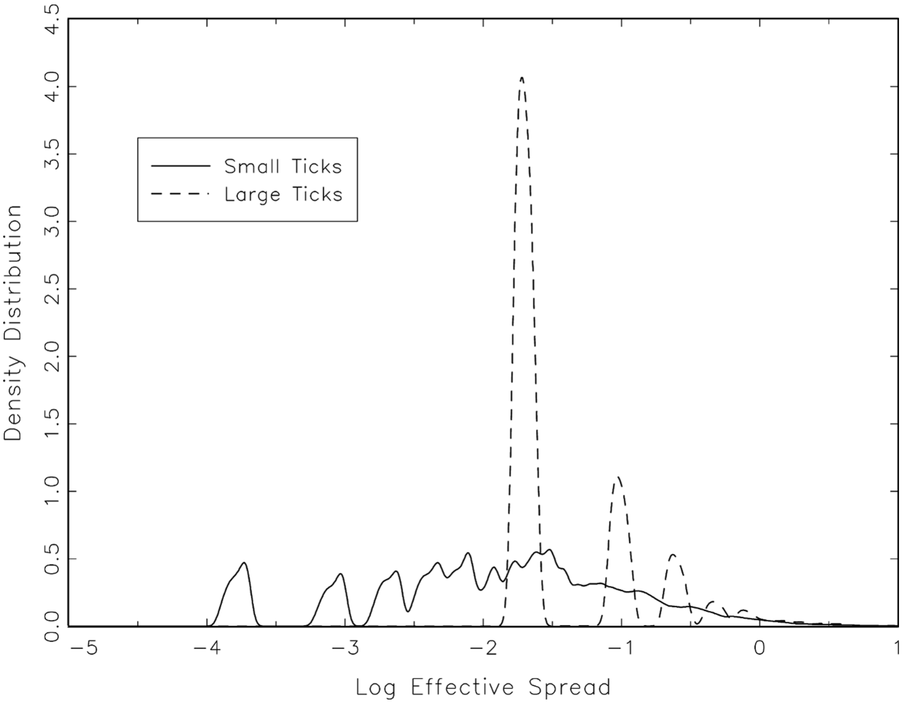

FIGURE 3. The kernel density estimate of the log effective spread is presented for 2,540,764 trades in the small tick regime and 2,155,658 trades in the large tick regime.

regime has a significantly higher log range for every percentile except the 1% percentile. Hourly range measures with no price movements are more frequent in the large tick regime than in the small tick regime. Apart from this irregularity at the first percentile, however, we find that the volatility is strictly higher for large ticks. Figure 5 shows the cumulative distribution of the log range by tick size regime. The large tick cumulative distribution is strictly to the right of the small tick cumulative distribution above the 2% quantile. We conclude that larger ticks increase price volatility measured by the log range. The median increase in the percentage hourly range amounts to approximately 20% from $e^{-0.5667} - 0.1 = 0.4674$ to $e^{-0.4160} - 0.1 = 0.5597$. This represents an economically significant increase which is of the same magnitude as the transaction cost increase.

Next, we confirm these findings with a formal panel regression analysis in Table 5. We use a Feasible General Least Square (FGLS) estimator which allows for panel specific serial correlation of the error. The most parsimonious specification (specification I) regresses the hourly volatility measures (LRANGE) on a regime dummy for large ticks (TICK DUMMY), fixed effects for each intraday trading hour (not reported) and fixed effects for each stock (not reported). The fixed effects for each trading hour control for the intraday patter of stock

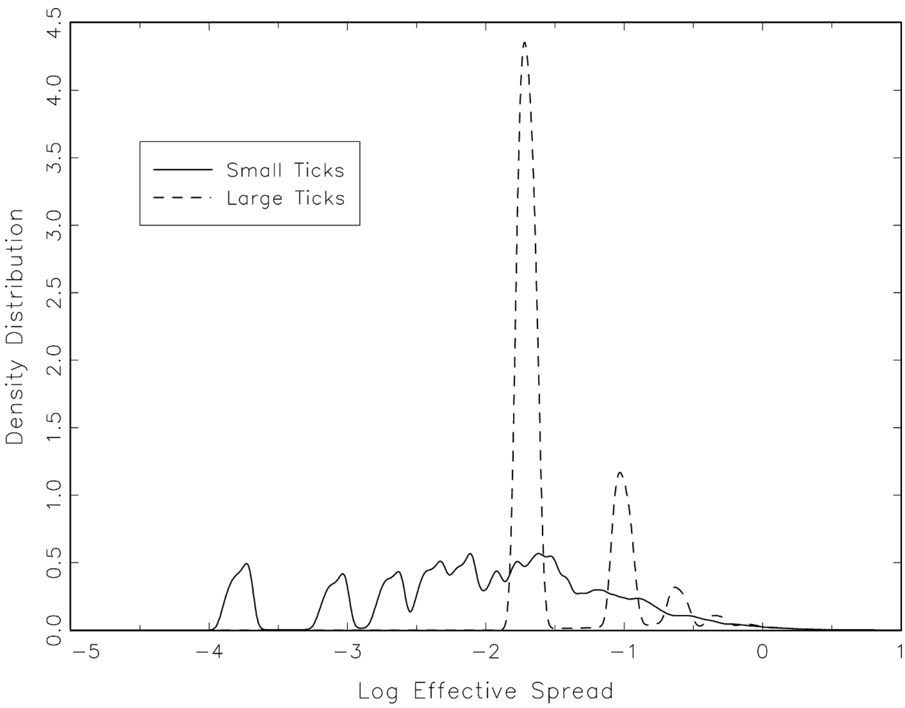

FIGURE 4. The kernel density estimate of the log effective spread is presented for 1,587,683 executed orders in the small tick regime and 1,331,146 executed orders in the large tick regime.

price volatility. Fixed effects for each stock can be included because 21 of the 26 stocks trade sequentially in both tick size regimes. Hence we control for heterogenous stock specific volatility levels and the regime dummy captures only the average volatility change of each stock when it moves from one tick size regime to the other. The regime dummy is statistically highly significant. A regression coefficient of 0.25 for the TICK DUMMY implies that the volatility measure 0.1+ RANGE increases by a factor of $e^{0.25} = 1.28$, or 28%, across tick size regimes. The average hourly LRANGE is approximately $-0.54 = \ln(0.1+0.48)$ under small ticks. The RANGE itself therefore increases by a factor Δ, where $-0.54 + 0.25 = \ln(0.1 + 0.48 \times \Delta)$. The corresponding point estimate for the increase in the percentage hourly range follows as $\Delta - 1 = 0.35$, or 35%. Hence, the volatility increase is also significant in economic terms.

The baseline specification does not account for common volatility effects across stocks. Regression specification II includes current and lagged hourly index volatility (LRANGE(INDEX)) as independent variables. Thus, we capture the common volatility component across stocks. Hourly index volatility is measured by the log range of the CAC40 index over every full trading hour just like individual stock volatility. Controlling for marketwide volatility effects, we obtain

TABLE 3. Summary statistics on hourly log percentage range by stock.

	Small ticks					Large ticks				
Stock Name	Obs.	Mean	St. dev.	Skew.	Kurt.	Obs.	Mean	St. dev.	Skew.	Kurt.
Sanofi	609	−0.679	0.627	0.245	2.945	1,460	−0.226	0.601	−0.091	3.134
Total	1,033	−0.617	0.545	0.072	2.598	582	0.012	0.513	−0.126	2.926
Elf Aquitaine	646	−0.765	0.518	0.135	2.703	630	−0.170	0.526	−0.068	2.895
Bouygues	753	−0.756	0.649	0.028	2.805	3,785	−0.582	0.659	−0.330	3.403
Lyonnaise des Eaux	2,886	−0.761	0.595	0.150	2.812	1,441	−0.540	0.618	−0.243	3.568
Lafarge	945	−0.263	0.613	0.083	2.555	627	0.010	0.547	−0.258	3.385
AXA	615	−0.569	0.576	0.812	4.580	269	0.046	0.565	−0.230	3.687
CFF	52	−0.771	0.683	0.277	2.442	336	−0.437	0.749	−0.199	3.208
BIC	1,230	−0.247	0.634	0.086	2.733	96	−0.021	0.601	0.345	2.573
Bancaire (CIE)	629	−0.588	0.655	−0.145	2.649	2,661	−0.521	0.700	−0.338	3.192
Général des Eaux	1,091	−0.545	0.614	0.113	2.895	2,016	−0.571	0.570	−0.180	3.385
Spie Batignolles	404	−0.281	0.606	0.364	2.860	58	−0.035	0.616	−0.055	3.071
Paribas	1,165	−0.304	0.603	0.136	2.777	461	−0.125	0.552	−0.185	4.262
CCF	869	−0.262	0.667	0.151	2.729	288	−0.098	0.590	−0.121	3.277
Alcatel Alstom	3,201	−0.626	0.603	0.335	3.143	486	−0.197	0.741	0.474	3.594
Havas	2,291	−0.630	0.628	0.090	2.813	275	−0.232	0.558	−0.253	4.488
Valéo	741	−0.071	0.671	0.023	2.959	457	0.043	0.546	0.048	3.051
Société Général	110	−0.392	0.521	−0.093	3.107	3,255	−0.583	0.594	0.051	3.777
Crédit LCL France	2,758	−0.816	0.629	0.168	2.852	1,256	−0.366	0.607	−0.267	3.556
BNP	396	−0.039	0.602	0.028	2.576	363	−0.029	0.568	0.119	3.989
SGS Thomson	648	−0.130	0.668	−0.129	2.580	115	−0.071	0.595	−0.832	5.320
France Télécom	409	0.191	0.562	−0.134	3.020	—	—	—	—	—
Accor	—	—	—	—	—	799	−0.611	0.607	−0.236	3.523
Peugeot	—	—	—	—	—	1,181	−0.429	0.654	−0.196	3.180
Saint-Gobain	—	—	—	—	—	711	−0.432	0.535	−0.210	3.824
Canal Plus	—	—	—	—	—	124	−0.422	0.617	−0.199	2.679
All Stocks	23,481	−0.542	0.661	0.186	2.881	23,732	−0.423	0.651	−0.190	3.456

Notes: Distribution statistics of the hourly log percentage range (LRANGE) are calculated for all stocks in the French CAC40 index which trade for at least 20 hours in the price range (in French francs, FF) from FF 400 to FF 600 over the four-year period from January 1995 to December 1998. The log percentage range over hourly trading periods I_t from 10:00 a.m. to 5:00 p.m. is defined as

$$\mathrm{LRANGE}_t = \ln\left[0.1 + 200 \times \frac{\max_{\{s\in I_t\}} P_s^M - \min_{\{s\in I_t\}} P_s^M}{\max_{\{s\in I_t\}} P_s^M + \min_{\{s\in I_t\}} P_s^M}\right],$$

where P_s^M denotes the midprice between the best bid and ask quotes. The range measures are analyzed according to the tick size regime to which each stock is subject. Stocks quoted in the price range from FF 400 to FF 500 are subject to a minimum tick size of FF 0.1 (Small ticks), and stocks in the price range from FF 500 to FF 600 are subject to a minimum tick size of FF 1 (Large ticks).

TABLE 4. Distribution of hourly log percentage range by tick regime.

	Small ticks			Large ticks		
Percentiles	Centiles	1% C.I.		Centiles	1% C.I.	
1%	−1.9623	−1.9640	−1.9614	−2.3026	−2.3026	−2.3026
5%	−1.5838	−1.5841	−1.5832	−1.3102	−1.3102	−1.3102
10%	−1.3744	−1.3746	−1.3743	−1.2677	−1.2677	−1.2677
25%	−1.0087	−1.0092	−1.0085	−0.7959	−0.7959	−0.7959
50%	−0.5667	−0.5667	−0.5667	−0.4160	−0.4160	−0.4160
75%	−0.1038	−0.1043	−0.1035	0.0090	0.0090	0.0090
90%	0.3395	0.3393	0.3401	0.3870	0.3864	0.3870
95%	0.5925	0.5921	0.5925	0.6169	0.6164	0.6172
99%	1.0453	1.0448	1.0468	1.0696	1.0659	1.0715
Observations		23,481			23,732	
Mean			−0.5421			−0.4229
Std. dev.			0.6612			0.6509
Skewness			0.1860			−0.1902
Kurtosis			2.8814			3.4560

Notes: Hourly price volatility is calculated as the log percentage range between the highest and lowest quoted midprice over 60-minute intervals I_t from 10:00 a.m. to 5:00 p.m. for all stocks in the French CAC40 index which trade for at least 20 hours in the price interval between FF 400 and FF 600 over the four-year period from January 1995 to December 1998. For a midprice P_s^M given as the arithmethic average between the best bid and ask price, we define the log percentage range as

$$\text{LRANGE}_t = \ln\left[0.1 + 200 \times \frac{\max_{\{s\in I_t\}} P_s^M - \min_{\{s\in I_t\}} P_s^M}{\max_{\{s\in I_t\}} P_s^M + \min_{\{s\in I_t\}} P_s^M}\right].$$

The distribution of the hourly log percentage range is provided according to the tick size regime to which each stock is subject. Stocks quoted in the price range from FF 400 to FF 500 are subject to a minimum tick size of FF 0.1 (Small ticks), and stocks in the price range from FF 500 to FF 600 are subject to a minimum tick size of FF 1 (Large ticks).

a substantial increase in the log likelihood function. The TICK DUMMY is approximately 0.23 and therefore only slightly lower than in the baseline specification. This results does not change as we vary the number of lags for the index volatility included in the regression.

Specification III also allows for serial correlation of the individual stock volatility by inclusion of lagged dependent variables (LRANGE(−n)). Lagged stock volatility is highly significant and further improves the regression fit. The positive regime effect remains statistically highly significant. The TICK DUMMY coefficient drops to 0.1241. This is not surprising. Inclusion of lagged dependent variables in the specification implies that the TICK DUMMY coefficient captures only the short-run effect of the regime change. We can recover the permanent tick size effect by rescaling the coefficient by the factor $1/(1 - \sum_{i=1}^{7} \beta_i)$, where β_i represents the coefficient on the lagged dependent variable STOCKVOL with lag i. The long-run volatility effect of large ticks follows as 0.2293. This point estimate corresponds to a 31% increase in the percentage hourly range (RANGE).

We perform a variety of robustness checks on these regression results. Excluding 1% outliers in each tail of the volatility distribution did not qualitatively effect the results. Similarly, quantile regressions also produce a highly significant positive regime dummy effect under each of the above regression specifications. We

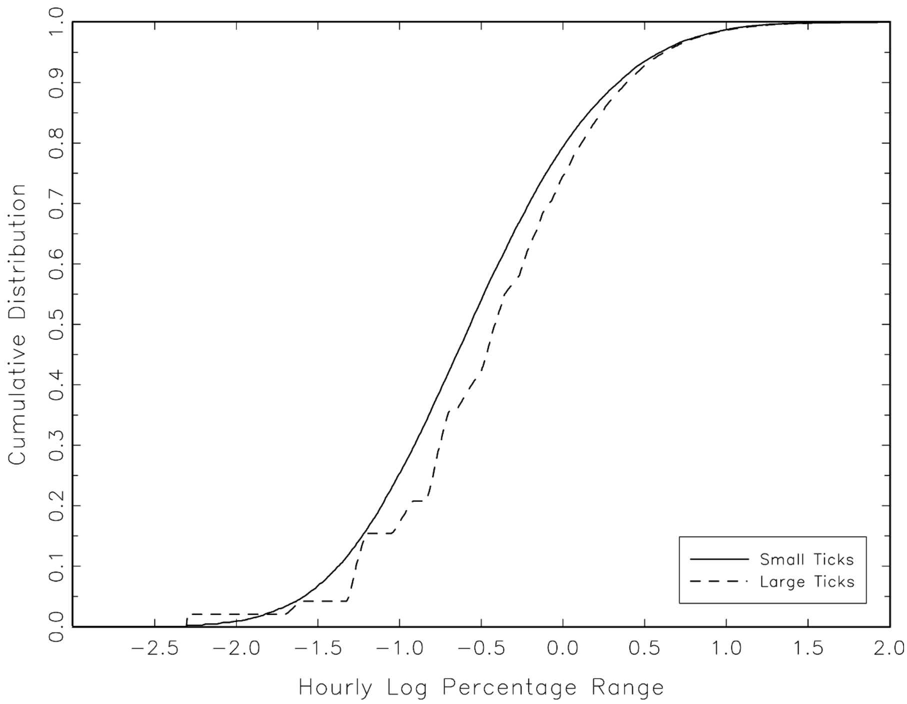

FIGURE 5. The cumulative distribution of the hourly log percentage range for 23,481 volatility measures in the small tick regime and 23,732 volatility measures in the large tick regime.

can therefore assert that the results are not induced by a relatively small number of volatility outliers.

7. Conclusion

We analyze the causal linkage between transaction costs and financial price volatility using a cross-sectional identification of the transaction cost differences based on exogenous tick size regulation related to the stock price level. It is shown that an increase in the tick size at FF 500 in the French stock market increases the median effective spread, and therefore transaction costs, by approximately 20% for stocks in the CAC40 index. This finding corresponds to qualitatively similar results in the existing literature. In a second step, we use these exogenous transaction cost differences to explore the volatility implication. Here the paper makes two important contributions to the existing literature. First, the panel data structure allows us to identify and separately control for time changing volatility unlike the data structure available from one-time marketwide regulatory tick size changes. Second, we avoid the pitfalls of biased volatility measurement common

TABLE 5. Volatility regressions.

	LRANGE (= log hourly percentage range of the stock midprice)		
Dependent Variable	Specification I	Specification II	Specification III
TICK DUMMY	0.2535***	0.2258***	0.1241***
(Large ticks = 1)	(0.0091)	(0.0074)	(0.0059)
LRANGE(INDEX)	—	0.5680***	0.5479***
		(0.0061)	(0.0061)
LRANGE(INDEX)(−1)	—	0.0928***	−0.0441***
		(0.0060)	(0.0068)
LRANGE(INDEX)(−2)	—	0.0565***	−0.0277***
		(0.0060)	(0.0068)
LRANGE(INDEX)(−3)	—	0.0344***	−0.0281***
		(0.0061)	(0.0067)
LRANGE(−1)	—	—	0.2070***
			(0.0047)
LRANGE(−2)	—	—	0.0805***
			(0.0048)
LRANGE(−3)	—	—	0.0499***
			(0.0048)
LRANGE(−4)	—	—	0.0346***
			(0.0045)
LRANGE(−5)	—	—	0.0233***
			(0.0045)
LRANGE(−6)	—	—	0.0276***
			(0.0045)
LRANGE(−7)	—	—	0.0358***
			(0.0043)
HOUR DUMMIES	Yes	Yes	Yes
STOCK DUMMIES	Yes	Yes	Yes
Observations	47, 213	45, 950	44, 731
Stocks	26	26	26
Wald $\chi^2(n = 32, 36, 43)$	12940***	27436***	41708***
Log Likelihood	−36680	−31167	−29799
Permanent effect			
TICK DUMMY			0.2293

Notes: Stock price volatility, measured by the hourly log percentage range of the quoted midprice (LRANGE), is calculated for all stocks in the French CAC40 index which trade for at least 20 hours in the price interval between FF 400 and FF 600 over the four-year period from January 1995 to December 1998 and regressed on a dummy variable of the tick size regime (TICK DUMMY) as an exogenous transaction cost proxy. Stocks with "low transaction costs" in the price range from FF 400 to FF 500 are subject to a minimum tick size of FF 0.1 (Small ticks), and stocks with "high transaction costs" in the price range from FF 500 to FF 600 are subject to a minimum tick size of FF 1 (Large ticks). We use an FGLS estimator which allows for serial error correlation specific to each panel. LRANGE(INDEX) measures the hourly log percentage range of the CAC40 index and LRANGE(INDEX)(−1) the corresponding range lagged by one trading hour. LRANGE(−1) denotes the lagged (by 1 trading hour) hourly log percentage range of the individual stock. Not reported are additional intraday dummies for each full trading hour (HOUR DUMMIES) and fixed effects for each stock (STOCK DUMMIES). Standard errors are provided in parentheses, and significance levels at 5% (*), 3% (**), and 1% (***) level are marked.

in previous studies based on the standard deviation or variance of midprices. Instead, we use the so-called percentage range as a volatility metric which is unbiased with respect to price measurement under different tick size regimes. Our inference is therefore both statistically powerful and unbiased.

Our data sample yields 47,213 hourly range measures for all stocks in the CAC40 index trading around the tick size discontinuity at FF 500. Panel regressions with stock specific fixed effects show a statistically strong effect of the exogenous tick regime dummy on individual stock volatility even after controlling for marketwide volatility. The volatility increase attributed to the tick regime dummy and measured by the percentage hourly range is more than 30%, and therefore of a similar magnitude as the tick size induced transaction cost increase.

We therefore conclude that the effect of transaction costs on volatility is positive and significant, both statistically and economically. The general volatility increase registered for U.S. stock markets (Campbell et al. 2001) is therefore unlikely to be explained by the important transaction cost decrease in the same markets over the last two decades. A more competitive tick size structure with lower feasible minimum price variations is, on the contrary, likely to reduce financial price volatility. On the policy side, security transaction costs should increase rather than decrease return volatility. Our volatility measures are likely to underestimate the destabilizing role of security transaction taxes because they—unlike large ticks—also reduce the stabilizing liquidity supply. In the light of our evidence and the liquidity supply argument, a security transaction tax should be deemed counterproductive. On the larger issue of short-term speculation and financial price stability, our evidence supports Friedman's (1953) general defense of financial speculation. High transaction costs discourage short-term speculation, and this can explain why volatility increases whenever transaction costs increase.

References

Ahn, Hee-Joon, Charles Q. Cao, and Hyuk Choe (1996). "Tick Size, Spread and Volume." *Journal of Financial Intermediation*, 5, 2–22.

Ahn, Hee-Joon, Charles Q. Cao, and Hyuk Choe (1998). "Decimalization and Competition among Stock Markets: Evidence from the Toronto Stock Exchange Cross-listed Securities." *Journal of Financial Markets*, 1, 51–87.

Alizadeh, Sassan, Michal W. Brandt, and Francis X. Diebold (2002). "Range-Based Estimation of Stochastic Volatility Models." *Journal of Finance*, 57, 1047–1092.

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Heiko Ebens (2001). "The Distribution of Stock Return Volatility." *Journal of Financial Economics*, 61, 43–76.

Angel, James J. (1997). "Tick Size, Share Prices, and Stock Splits." *Journal of Finance*, 52(2), 655–681.

Bacidore, Jeffrey M. (1997). "The Impact of Decimalization on Market Quality: An Empirical Investigation of the Toronto Stock Exchange." *Journal of Financial Intermediation* 6, 92–120.

Bessembinder, Hendrik (2000). "Tick Size, Spreads, and Liquidity: An Analysis of Nasdaq Securities Trading Near Ten Dollars." *Journal of Financial Intermediation*, 9(3), 213–239.
Bessembinder, Hendrik (2003). "Trade Execution Costs and Market Quality after Decimalization." *Journal of Financial and Quantitative Analysis*, 38(4), 747–777.
Bessembinder, Hendrik, and Subhrendu Rath (2002). "Trading Costs and Return Volatility: Evidence from Exchange Listings." Working Paper, University of Utah.
Biais, Bruno, Pierre Hillion, and Chester Spatt (1995). "An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse." *Journal of Finance*, 50, 1655–1689.
Biais, Bruno, Pierre Hillion, and Chester Spatt (1999). "Price Discovery and Learning during the Preopening Period in the Paris Bourse." *Journal of Political Economy*, 107(6), 1218–1248.
Bollen, Nicolas P. B., and Robert E. Whaley (1998). "Are 'teenies' better?" *Journal of Portfolio Management*, 25(1), 10–26.
Bourghelle, David, and Fany Declerk (2004). "Why Markets Should not Necessarily Reduce Tick Size." *Journal of Banking and Finance*, Special Issue, 28(2), 373–398.
Campbell, John Y., and Kenneth Froot (1994). "International Experience with Security Transaction Taxes." In *The Internationalization of Equity Markets*, edited by Jeffrey Frankel. University of Chicago Press.
Campbell, John Y., Martin Lettau, Burton G. Malkiel, and Yexiao Xu (2001). "Have Individual Stocks become More Volatile? An Empirical Exploration of Idiosyncratic Risk." *Journal of Finance*, 56(1), 1–43.
De Jong, Frank, Theo Nijman, and Ailsa Röell (1995). "A comparison of the Cost of Trading French Shares on the Paris Bourse and on SEAQ International." *European Economic Review*, 39, 1277–1301.
Domowitz, Ian, Jack Glen, and Ananth Madhavan (2001). "Liquidity, Volatility and Equity Trading Costs across Countries and over Time." *International Finance*, 4(2), 221–255.
Dooley, Michael P. (1996). "Tobin Tax: Good Theory, Weak Evidence, Questionable Policy." In *The Tobin Tax: Coping with Financial Volatility*, edited by Mahub ul Haq, Isabelle Kaul, and Inge Grunberg. Oxford University Press.
Eichengreen, Barry, James Tobin, and Charles Wyplosz (1995). "Two Cases for Sand in the Wheels of International Finance." *Economic Journal*, 105, 162–172.
Friedman, Milton (1953). *Essays in Positive Economics*. University of Chicago, Press.
Goldstein, Michael A., and Kenneth A. Kavajecz (2000). "Eighths, sixteenths, and Market Depth: Changes in the Tick Size and Liquidity Provision on the NYSE." *Journal of Financial Economics*, 56, 125–149.
Harris, Laurence (1994). "Minimum Price Variations, Discrete Bid-ask Spreads, and Quotation Size." *Review of Financial Studies*, 7, 149–178.
Harris, Laurence (1996). "Does a Large Minimum Price Variation Encourage Order Exposure?" New York Stock Exchange Working Paper No. 96–05.
Jones, Charles M. (2002). A Century of Stock Market Liquidity and Trading Costs. Working Paper, Columbia University.
Jones, Charles M., and Marc L. Lipson (1999). "Execution Costs of Institutional Equity Orders." *Journal of Financial Intermediation*, 8, 123–140.
Jones, Charles M., and Marc L. Lipson (2000). "Sixteenths: Direct Evidence on Institutional Trading Costs." *Journal of Financial Economics*, 59(2), 253–278.
Jones, Charles M., and Paul J. Seguin (1997). "Transaction Costs and Price Volatility: Evidence from Commission Deregulation." *American Economic Review* 87(4), 728–737.
Keynes, John M. (1936). *The General Theory of Employment, Interest and Money*. McMillan Press.
Lau, Sie Ting, and Thomas H. McInish (1995). "Reducing Tick Size on the Stock Exchange of Singapore." *Pacific Basin Finance Journal*, 3, 485–496.
Miller, Merton H. (1991). *Financial Innovations and Market Volatility*, Basil Blackwell.

Mulherin, Harold J. (1990). "Regulation, Trade Volume, and Stock Market Volatility." *Revue Economique*, 41(5), 923–937.

Porter, David C., and Daniel G. Weaver (1997). "Tick Size and Market Quality." *Financial Management*, 26, 5–26.

Ronen, Tavy, and Daniel G. Weaver (2001). "Teenies' Anyone?" *Journal of Financial Markets*, 4, 231–260.

Schwert, G. William, and Paul J. Seguin (1993). "Security Transaction Taxes: An Overview of Costs and Benefits and Unresolved Questions." *Financial Analysts Journal*, 49(5), 27–35.

Stiglitz, Joseph E. (1989). "Using Tax Policy to Curb Speculative Short-term Trading." *Journal of Financial Service Research*, 3(2–3), 101–115.

Summers, Laurence H., and Victoria P. Summers (1989). "When Financial Markets Work Too Well: A Cautious Case for a Security Transaction Tax." *Journal of Financial Service Research*, 3(2–3), 261–286.

Tobin, James E. (1978). "A Proposal for International Monetary Reform." *Eastern Economic Journal*, 4, 153–159.

Tobin, James E. (1984). "On the Efficiency of the Financial System." *Lloyds Bank Review*, 153 (July), 1–15.

ul Haq, Mahbub, Isabelle Kaul, and Inge Grunberg, eds. (1996). *The Tobin Tax: Coping with Financial Volatility*. Oxford University Press.

Umlauf, Steven R. (1993). "Transaction Taxes and the Behavior of the Swedish Stock Market." *Journal of Financial Economics*, 33, 227–240.

Venkataraman, Kumar (2001). "Automated versus Floor Trading: An Analysis of Execution Costs on the Paris and New York Stock Exchange." *Journal of Finance*, 56(4), 1445–1485.