In Search of International Integration:

An Examination of Intraday North American Trading

of Canadian Dually Listed Stocks

Aditya Kaul[†]
Faculty of Business
University of Alberta
akaul@gpu.srv.ualberta.ca
Tel. 780.492.5027
Fax. 780.492.3325

Vikas Mehrotra
Faculty of Business
University of Alberta
vmehrotr@gpu.srv.ualberta.ca
Tel. 780.492.2976
Fax. 780.492.3325

This draft: *March 14, 2000*

[†] Corresponding author. The authors thank seminar participants at the 1999 Northern Finance Association meetings and at the Universities of Alberta and Oregon, and John Chalmers, Larry Dann, Jarrad Harford, Wayne Mikkelson, and especially Rich Lyons and Barry Scholnick for helpful suggestions, Kamal Khanna of the Ontario Securities Commission for regulatory information and Olsen and Associates for providing exchange rate data. We are responsible for any remaining errors.

In Search of International Integration: An Examination of Intraday North American Trading of Canadian Dually Listed Stocks

Abstract

We examine international equity market integration using intraday data for a sample of Canadian stocks that trade on both the Toronto Stock Exchange (TSE) and the New York Stock Exchange (NYSE). There are several advantages of this sample: Canadian stocks trade as stocks (not ADRs) in the U.S.; investors in each country are free to route trades to the foreign market; the additional cost of trading abroad is small, especially for institutional traders; and the markets have perfectly synchronous trading hours. These conditions allow us to examine both whether the *law of one price* holds and the process by which deviations from it are corrected. Our tests show the U.S. dollar prices of the stocks in the two markets to be cointegrated. The estimated error correction models show rapid adjustments to deviations from the law of one price, with more aggressive adjustment generally occurring on the thinner market. The ratio of the NYSE price to the TSE price (an estimate of the spot exchange rate) is tightly distributed around the actual spot rate. The frequency of arbitrage opportunities is low, and the profits seldom exceed reasonable estimates of trading costs. In general, trading for more actively traded stocks exhibits stronger integration. Puzzlingly, NYSE order flow alone responds significantly to price deviations. This single market order flow response is at odds with the rest of our evidence, which favors integrated trading.

Introduction

The *law of one price* asserts that, in perfectly integrated global markets, identical goods and services trade at identical prices. If international financial markets are integrated, deviations from the law of one price will usually be small enough to preclude arbitrage; significant deviations should be corrected rapidly.

In this paper, we consider international market integration using intraday data for Canadian stocks that trade both on the Toronto Stock Exchange (the TSE) and on the New York Stock Exchange (the NYSE). There are several advantages of this sample. Canadian shares trade as shares (not ADRs) in both markets, and the incremental costs of trading on the foreign market are small, especially for institutional traders. Finally, both the TSE and the NYSE are open for business from 9:30 a.m. until 4:00 p.m. EST, allowing us to examine a full day of simultaneous trading in two international venues.

Our research makes two important contributions. The first of these is methodological. While we examine the extent to which the law of one price — a long–run equilibrium relation — holds, we recognize that it is unlikely to hold exactly at any instant. Hence, we treat market integration primarily as a disequilibrium process, and examine it in terms of the magnitude and duration of price deviations, the process by which these deviations are corrected, and the size and frequency of arbitrage opportunities. By contrast, most earlier work on international integration has focused either on long–period price parity for goods and services (the large literature on Purchasing Power Parity, or PPP) or on international asset pricing.[1]

Our second contribution is an explicit recognition of the importance of trading activity in achieving international integration. Some stocks in our sample trade actively on the NYSE, some trade actively on the TSE, and the rest display intermediate levels of trading activity. This allows us to view integration as a function of trading activity in each market. We expect the law of one price to hold to a greater degree for stocks that are heavily traded in both markets, because each market offers sufficient liquidity to facilitate price–equalizing trades. Additionally, the thinner market (be it the TSE or the NYSE) is expected to display greater adjustment.

Our tests use intraday data from the stock and the foreign exchange markets from January 1 through December 31, 1995. We start our examination with cointegration tests. The law of one price implies that the difference between TSE and NYSE prices (expressed in the same currency) never becomes large even though both price series are non–stationary. Furthermore, if the two price series are cointegrated, we are able to estimate error correction models. The coefficients in these models shed light on the speed with which deviations from the law of one price are eliminated, and the extent to which prices in each market respond to these deviations.

If prices in the two markets do adjust to deviations from the law of one price, it is of interest to examine whether this adjustment is related to order flow. If trading is integrated, orders will be routed to the market offering the better price. We therefore relate signed order flow in the two markets to the deviation from the law of one price. Integration predicts a significant order flow response in each market.

An alternative form of the law of one price implies that ratio of the NYSE price to the TSE price is an estimate of the spot exchange rate. Using intraday prices from the two markets, we are able to construct several contemporaneous estimates of the spot rate implied by the ratio of prices, one per stock. If the markets are integrated, the various estimates derived from the different stocks will be very similar to each other and to the actual spot rate. We therefore examine the dispersion and bias of the cross-sectional distribution of the implied exchange rate around the spot rate.

A question of practical relevance is the frequency with which deviations from the law of one price allow arbitrage opportunities. We incorporate the bid-ask spreads from the TSE and NYSE and the inter-

---

[1] For recent surveys of PPP, see Dornbusch (1987) and Rogoff (1996). Harvey (1991) is one of the many studies of

bank market to estimate the frequency and magnitude of arbitrage opportunities for our sample of stocks. In integrated markets, these opportunities should be infrequent and should not provide large returns.

Our main results are as follows. For each of the dually listed stocks in our sample, the NYSE and TSE prices (both expressed in U.S. dollars) are cointegrated. The estimated error correction models show that large deviations from the law of one price are corrected rapidly, in thirty minutes or less, with prices in the thinner market generally adjusting to a greater extent. The cross sectional distribution of the exchange rate implied by NYSE and TSE prices is tightly distributed around the spot rate, and with little evidence of bias. The frequency of arbitrage opportunities is low and the returns in most cases are unlikely to cover reasonable estimates of transactions costs. Conditioning on trading activity, we find that less actively traded stocks display greater dispersion in the implied exchange rate, but a surprisingly rapid adjustment to deviations from the law of one price. Finally, NYSE order flow responds strongly to deviations from the law of one price, but TSE order flow does not. The order flow evidence muddies an otherwise favorable picture regarding integration of trading in the two markets.

This paper is related to Kleidon and Werner (1996), who examine intraday patterns on the London Stock Exchange (LSE) and the NYSE for a sample of 23 cross listed British stocks. The elevated volume and volatility at the NYSE open are indicative of private information revelation through trading and suggest that informed U.S. traders are unwilling to trade in London, which opens six hours before the NYSE. As Kleidon and Werner note, however, an alternative explanation for their results is the lack of complete substitutability between the underlying stocks trading in London and the American Depository Receipts (ADRs) trading on the NYSE. We circumvent this difficulty by focusing on a sample of perfect substitutes. Moreover, due to the completely synchronous trading hours for the TSE and NYSE, we are able to examine the law of one price and characterize the adjustment process.

---

international asset pricing.

This paper is also related to research that examines international transmissions across stock markets. Several studies have studied the links between stocks that are listed in multiple locations, e.g. Neumark, Tinsley and Tosini (1991) and Lieberman, Ben-Zion and Hauser (1999). However, these studies have not directly considered either the law of one price or the adjustment mechanism.

Finally, the results in this paper are related to the PPP literature. Many papers have examined whether the law of one price (which forms the basis of PPP) holds in goods markets, and have found that deviations are large and long-lived.[2] Our evidence is more favorable, and highlights the influence of transactions costs and imperfect substitutes on the extant results in the PPP literature.[3]

The rest of this paper is organized as follows. Section 1 details the advantages of studying integration using Canadian dually listed stocks. Section 2 describes our tests. Section 3 discusses our data and empirical methods, and Section 4 contains our results. Conclusions appear in Section 5.

1 Why use Canadian dually listed stocks to study integration?

As noted in the previous section, there are several advantages to using Canadian dually listed stocks to study international integration. We discuss these advantages next.

First, unlike companies from most other foreign countries, Canadian companies can list their shares directly in the U.S. Unless firms from other countries are willing to adhere to strict reporting standards, their stocks can only trade in the form of American Depository Receipts (ADRs).[4] ADRs are securities issued by a bank, which holds the underlying shares, and are not directly convertible into shares in the home country. Kleidon and Werner (1996) are unable to separate the conclusion that trading of British

---

[2] See, for instance, Engel (1993) and Froot, Kim and Rogoff (1995).

[3] Our approach is distinct from the asset pricing approach to studying international market integration. Asset pricing tests of integration focus on whether assets in different jurisdictions obey the same pricing relation (see, for instance, Jorion and Schwartz, 1986, and Harvey, 1991). By contrast, we examine the process through which markets come to be integrated.

[4] Most notably, companies must furnish accounting statements that are in accordance with GAAP.

stocks in London and of their ADRs in New York is segmented from the alternative that the ADRs and underlying shares are imperfect substitutes. By contrast, the complete substitutability between shares in Canadian firms bought on the TSE and the NYSE makes it more likely that Canadian and U.S. investors will be indifferent as to choice of trading venue. Note that, even for our sample, indirect benefits (such as soft dollar arrangements with brokers) might cause investors to trade on the domestic market.

Second, the costs of trading domestically and internationally do not differ substantially, even for small traders. For example, a Canadian investor will pay brokerage fees of USD (U.S. dollar) 43 for a $2000 trade routed to the NYSE, compared to $43 if the trade is sent to the TSE. The additional cost of routing her trades to the NYSE amounts to one percent on a $2000 order, and declines as the size of the order increases. Furthermore, the costs of foreign trades are likely lower for larger institutional traders (due, for instance, to the possibility of trading desks in both countries). Thus, we do not believe that the incremental out-of-pocket trading costs drive the decision of where to trade. In fact, our discussions with Canadian money managers confirm that the practice of trading in the U.S. markets is widespread. Likewise, U.S. money managers claim that they often acquire Canadian shares on the TSE.

Third, trading is possible on the NYSE and the TSE over the same period each day, between 9:30 a.m. and 4:00 p.m. EST. As a result, an investor buying or selling Canadian dually listed shares between these hours can choose between two international venues. The overlap in trading hours allows us to examine both the incidence of deviations from the law of one price, and the adjustment in both markets to these deviations. Most other pairs of major international stock exchanges overlap incompletely. In the cases where markets do not overlap at all, it becomes necessary to analyze the relation between trading period prices (returns) in one market and non - trading period prices (returns) in another, or the lead-lag relation in cross–market returns.[5]

[5] See, for instance, Neumark et al. (1991) and Lieberman et al. (1999). Karolyi (1995) examines the return and volatility relation between the S&P 500 index and the TSE 300 index and portfolios of dually listed and non–dually listed Canadian stocks. However, Karolyi uses close-close returns and market aggregates.

2 Tests of integrated markets

If trading of Canadian dually listed stocks on the TSE and the NYSE is integrated, the stocks will obey the law of one price. This law states that a Canadian stock will have the same price at the same instant on the NYSE and the TSE, when the latter price (in \$) is converted into U.S. dollars:

$$P_{j,t}^{US} = P_{j,t}^{TSE} S_t(USD/\$) \qquad \langle 1 \rangle$$

where $P_{j,t}^{US}$ and $P_{j,t}^{TSE}$ are the prices of stock j on the NYSE and the TSE at time t and $S_t(USD/\$)$ is the spot price of the Canadian dollar at time t. For notational convenience the spot rate is henceforth expressed as S_t. Expression ⟨1⟩ says that prices of stock j should be affected by the same fundamentals but that the exchange rate should be instrumental in determining their relative position.

We have assumed that the spot rate is exogenous. This is reasonable, given the activity levels in the two markets. The daily volume of foreign exchange trading far exceeds trading volume on the major stock exchanges, and, as one of the actively traded currencies, the Canadian dollar experiences daily trading far exceeding the volume of trading in any dually listed stock.[6] Aggregate portfolio flows into or out of the Canadian equity market could conceivably affect the price of the Canadian dollar; however, such effects are likely to be of a long-horizon nature. Our horizon (a day or less) is short enough to make our assumption of exchange rate exogeneity reasonable.

In our analysis of North American trading of dually listed Canadian stocks we conduct four tests. Three of these tests are based on expression ⟨1⟩ and discussed in Sections 2.1, 2.3 and 2.4. The fourth test is an analysis of order flow and is described in Section 2.2.

---

[6] According to the most recent figures released by the Bank for International Settlements (April 1998), the daily value of trading in the Canadian dollar averaged USD29.80 billion (spot plus futures volume) or USD11.2 billion (spot alone). By contrast, the value of daily trading for all stocks on the TSE is approximately USD1 billion.

2.1 Cointegration tests

Unless the TSE and the NYSE are perfectly integrated and informationally efficient, we do not expect expression ⟨1⟩ to hold at every instant. In fact, it may not even be possible for two markets to be perfectly integrated if information is costly to transmit from one market to the other. However, we expect ⟨1⟩ to be a reasonable characterization of equilibrium price behavior, and do not expect deviations from ⟨1⟩ to be large or persistent. Cointegration tests allow us to examine both long-run price behavior and the process by which disequilibrium price deviations are corrected.[7]

The stock prices in ⟨1⟩ are likely to be non-stationary. However, Engle and Granger (1987) point out that two nonstationary series need not diverge. To the extent that prices on the NYSE and the TSE are subject to arbitrage or least-cost dealing considerations, the difference between the U.S. dollar prices in the two markets should never become indefinitely large, even if each price is nonstationary. This stationarity of the difference between prices in the two markets means that $P_{j,t}^{US}$ and $P_{j,t}^{TSE} S_t$ are cointegrated.

We test whether the U.S. dollar prices of the dually listed stocks are cointegrated. However, from our standpoint, perhaps of greater interest is Engle and Granger's insight that cointegrated series can be expressed in error correction model form.[8] Define the error in interval *t* as

$$m_{j,t} = P_{j,t}^{TSE} S_t - P_{j,t}^{US} \qquad \langle 2 \rangle$$

Then, according to the Granger Representation Theorem (see Engle and Granger (1987)), returns for stock *j* in the two markets can be written as:

[7] Harris et. al (1995) and Hasbrouck (1995) have applied cointegration to price discovery on the NYSE and satellite markets.

[8] Our cointegration tests are based on the natural logarithm of U.S. dollar prices in the two markets. The error correction model involves first differences of the cointegrated variables, and the use of logs conveniently allows us to express the model in terms of returns.

$$R_{j,t+1}^{TSE} = \alpha_0^{TSE} + \sum_{l=0}^{L_1} \alpha_{TSE,l}^{TSE} R_{j,t-l}^{TSE} + \sum_{l=0}^{L_2} \alpha_{US,l}^{TSE} R_{j,t-l}^{US} + \gamma_j^{TSE} \mu_{j,t} + \varepsilon_{j,t+1}^{TSE} \quad \langle 3a \rangle$$

$$R_{j,t+1}^{US} = \alpha_0^{US} + \sum_{l=0}^{L_1} \alpha_{TSE,l}^{US} R_{j,t-l}^{TSE} + \sum_{l=0}^{L_2} \alpha_{US,l}^{US} R_{j,t-l}^{US} + \gamma_j^{US} \mu_{j,t} + \varepsilon_{j,t+1}^{US} \quad \langle 3b \rangle$$

where $R_{j,t+1}^{TSE}$ and $R_{j,t+1}^{US}$ are TSE and NYSE returns in interval $t+1$ (the subscript j on the autoregressive coefficients is suppressed).

Equations ⟨3a⟩ and ⟨3b⟩ are vector autoregressions (VARs) for U.S. dollar returns in the two markets augmented with the previous period's error. In order to interpret the above two equations, suppose that $\mu_{j,t} > 0$, which means that, for stock j in interval t, TSE prices (expressed in U.S. dollar) are higher than NYSE prices. In order for this error to be corrected, TSE prices must fall and NYSE prices must rise in interval $t+1$, ceteris paribus. This will be reflected in negative TSE returns and positive NYSE returns for stock j in interval $t+1$. Thus, we expect to find γ_j^{TSE} to be negative and γ_j^{US} to be positive.[9]

The magnitudes of γ_j^{TSE} and γ_j^{US} reveal the extent to which prices on the NYSE and the TSE adjust to errors and hence the importance of the two markets in price discovery for stock j. For instance, if price adjustment on the TSE is more gradual than on the NYSE, γ_j^{TSE} will be small and γ_j^{US} large in absolute value. Additional information is supplied by the incremental R^2 after adding $\mu_{j,t}$ to the VARs. The incremental R^2 indicates the importance of the lagged error in explaining returns in the two markets. Finally, the coefficients from the VARs can be used to examine the speed of adjustment to deviations from the law of one price.

---

[9] It is possible for prices in both markets to rise (fall), with the increase (decline) for the TSE being less (greater) than that for the NYSE. On average, however, subsequent price changes in the two markets should be random, and error correction should dominate the price adjustment.

2.2 Order flow tests

The order flow tests are designed to complement the cointegration tests and provide an understanding of the process through which markets come to be integrated. If the error correction models reveal significant price adjustment in response to deviations from the law of one price, it is logical to then ask how prices in the two markets come to adjust. If trading decisions are made in an integrated North American market, orders initiated both in Canada and in the U.S. will accumulate on one side in the two markets, buys in the cheaper market and sells in the more expensive market. This persistence in net (signed) order flow tips the market makers off about potential mispricing and induces price adjustments. Alternatively, it is possible that adjustment occurs in the absence of order flow. To the extent that market makers in each market watch prices in the other market and the exchange rate, they are likely to detect deviations quickly and respond, even before orders have arrived to capitalize on the difference in prices.

The two possibilities represent different statements about the behavior of order flow. The first predicts that NYSE and TSE order flow is sensitive to deviations from the law of one price. The second predicts that there is no necessary relation between order flow and deviations from the law of one price. We examine the link between order flow and deviations from the law of one price by estimating the following two vector autoregressions for signed order flow in the two markets, augmented with the previous period's error:

$$OF_{j,t+1}^{TSE} = b_0^{TSE} + \sum_{l=0}^{L_1} b_{TSE,l}^{TSE} OF_{j,t-l}^{TSE} + \sum_{l=0}^{L_2} b_{US,l}^{TSE} OF_{j,t-l}^{US} + y_j^{TSE} m_{j,t} + u_{j,t+1}^{TSE} \qquad \langle 4a \rangle$$

$$OF_{j,t+1}^{US} = b_0^{US} + \sum_{l=0}^{L_1} b_{TSE,l}^{US} OF_{j,t-l}^{TSE} + \sum_{l=0}^{L_2} b_{US,l}^{US} OF_{j,t-l}^{US} + y_j^{US} m_{j,t} + u_{j,t+1}^{US} \qquad \langle 4b \rangle$$

where $OF_{j,t+1}^{TSE}$ and $OF_{j,t+1}^{US}$ are signed order flow on the TSE and the NYSE in interval *t+1*. We use the number of transactions as our measure of order flow, since it is less sensitive to outliers than is volume. In order to interpret ⟨4a⟩ and ⟨4b⟩ suppose, once again, that $m_{j,t} > 0$ (TSE prices are higher than NYSE prices in period *t*). If order flow is sensitive to the deviation, and trading is integrated, there should be an imbalance of sells on the TSE and an imbalance of buys on the NYSE in interval *t+1*, ceteris paribus. Assigning buys a positive sign and sells a negative sign, we expect to find y_j^{TSE} to be negative and y_j^{US} to be positive. The magnitudes of y_j^{TSE} and y_j^{US} reveal the extent to which order flow on the TSE and the NYSE adjusts to errors. For instance, if the order flow response is sluggish on the TSE relative to that on the NYSE, y_j^{TSE} will be smaller than y_j^{US} in absolute value. We also provide the incremental R^2 after adding $m_{j,t}$ to the order flow VAR—as a measure of the importance of the price deviation in explaining order flow in the two markets.

2.3 Implied exchange rate tests

Section 2.1 has described a time series test of market integration. This test says is not informative about the magnitude of deviations from the law of one price. Several alternative tests can be derived from ⟨1⟩ to address this deficiency. One appealing test is based on the *implied exchange rate*. From ⟨1⟩ the ratio of the NYSE share price to the TSE price provides an estimate of the spot exchange rate:

$$S_{j,t}^{I} = \frac{P_{j,t}^{US}}{P_{j,t}^{TSE}} \qquad \langle 5 \rangle$$

where $S_{j,t}^{I}$ is the exchange rate implied by the prices of stock *j* at time *t*. Using prices of the dually listed stocks we can, at any instant, compute several estimates of the implied exchange rate, one per stock. In integrated markets, the different exchange rate estimates should be very similar and tightly clustered around the true spot exchange rate. Accordingly, we examine the dispersion of the distribution of the implied exchange rate around the spot rate. This dispersion should be small. Additionally, the distribution should be centered on the actual spot rate. Hence, we also examine whether the mean of the distribution is an unbiased predictor of the spot rate.

The implied exchange rate is the average of the bid and ask prices implied by stock market quotes. We start with the fact that converting U.S. dollars into $ (*buying* $) can be synthesized by buying shares on the NYSE and simultaneously selling shares on the TSE. This transaction is completed by using the shares bought in the U.S. to cover the short position in Canada. Likewise, converting $ into U.S. dollars (*selling* $) can be accomplished by selling shares on the NYSE and buying shares on the TSE. The first case yields the ask price of the $ and the second the bid price of the $. It is obvious, then, that $A^{I}_{j,t}$ and $B^{I}_{j,t}$, the ask and bid prices of the $ implied by the prices of stock *j*, are given by

$$A^{I}_{j,t} = \frac{A^{US}_{j,t}}{B^{TSE}_{j,t}} \text{ and } B^{I}_{j,t} = \frac{B^{US}_{j,t}}{A^{TSE}_{j,t}} \qquad \langle 6 \rangle$$

where $A^{US}_{j,t}$ and $A^{TSE}_{j,t}$ are the ask prices of stock *j* on the NYSE and on the TSE, and $B^{US}_{j,t}$ and $B^{TSE}_{j,t}$ are the corresponding bid prices. The implied exchange rate that we use in our tests is $S^{I}_{j,t} = 0.5(A^{I}_{j,t} + B^{I}_{j,t})$. The bias in the implied exchange rate for any interval is the difference between the mean implied exchange rate and the actual spot rate for that interval. The cross-sectional dispersion is computed by interval around the actual spot rate.

2.4 Arbitrage opportunities

A time–series version of the implied exchange rate test from the preceding section is informative about the prevalence of arbitrage opportunities as well as the size of such profits. The continuously compounded return on a simultaneous long position on the TSE and a short position on the NYSE for stock *j* is

$$ð^{1}_{j,t} = ln\left[\frac{B^{US}_{j,t}}{A^{TSE}_{j,t}\ S^{a}_{t}}\right] = ln\left[\frac{B^{I}_{j,t}}{S^{a}_{t}}\right] \qquad \langle 7a \rangle$$

where S_t^a and $B_{j,t}^I$ are the actual ask and implied bid exchange rates at time *t*. Similarly, the return on a short position on the TSE and a long position on the NYSE is

$$p_{j,t}^2 = ln\left[\frac{B_{j,t}^{TSE}\ S_t^b}{A_{j,t}^{US}}\right] = ln\left[\frac{S_t^b}{A_{j,t}^I}\right] \qquad \langle 7b\rangle$$

where S_t^b and $A_{j,t}^I$ are the actual bid and implied ask exchange rates at time *t*. We examine how frequently trades at the prevailing bid and ask prices generate positive $p_{j,t}^1$ or $p_{j,t}^2$.

It is also important to examine the size of the arbitrage profits. In order to simplify the computations we assume that all transactions occur at the midpoint of the bid and ask prices. In this case, ⟨7a⟩ and ⟨7b⟩ reduce to

$$p_{j,t}^1 = ln\left[\frac{S_{j,t}^I}{S_t}\right] = -p_{j,t}^2 \qquad \langle 8\rangle$$

where $S_{j,t}^I$ and S_t are the midpoints of the implied and actual exchange rates. Deviations between the implied and actual exchange rates represent the rates of return to arbitrage positions. Note that the two tests are complementary. In integrated markets, the frequency of deviations described by ⟨7a⟩ and ⟨7b⟩ should be small. Also, the returns in ⟨8⟩ should not be large.

3. Sample and data

We use data from one calendar year, 1995. We start with a sample of 177 stocks that were cross-listed on the TSE and a major U.S. exchange at the end of 1995. Stocks are excluded from further consideration if the median number of transactions per day per market was less than 1 in 1995; this reduces our sample to 151 stocks. Of these, 68 were listed on the NYSE and 83 traded on NASDAQ. We examine only NYSE listed stocks for two reasons. First the TSE by-laws afford best price execution protection only to prices from a list of recognized "auction" type exchanges (subject to minimum trading requirements). The NYSE and AMEX are part of this list, but NASDAQ is not. Second, trading activity tends to be over-stated on NASDAQ due to double counting of inter-dealer trades. We obtain intraday quote and transaction data from the TAQ and TSE tapes for 1995. Our tests also require exchange rate data. We obtain spot rate data for the Canadian dollar from Olsen and Associates. The dataset contains intraday time-stamped bid and ask quotes from major global banks, and has been used in several studies (see, for instance, Ito, Lyons and Melvin, 1998).

We retain the most recent quotes from the equity and foreign exchange markets, starting with the first quote of the day. This ensures that stale quotes from the previous day are excluded. . When returns are required, as in ⟨3a⟩ and ⟨3b⟩, the bid and ask prices from the stock market are used to compute midpoint returns.[10] Signed order flow is calculated using a variant of the Lee and Ready (1991) algorithm to sign trades. Specifically, the transaction price is compared with the midpoint of the bid and ask prices prevailing at the time of the transaction, so long as the quote has a time stamp at least five seconds before the time stamp of the transaction. (Otherwise the most recent quote satisfying the five second requirement is used.) A buy is a trade occurring at a price above the midpoint of the matching quote and a sell a trade occurring below the midpoint. Buys are assigned a positive sign and sells are assigned a negative sign.

One contribution of this paper is to examine integration as a function of trading activity. Accordingly, we sort our sample of dually listed stocks into three categories, corresponding to active, moderate, or light trading. We carry out this classification for trading activity on the TSE and on the NYSE, producing a 3-by-3 grid of nine cells. The (1,1) position in the grid consists of stocks that are lightly traded in both markets, while the (3,3) position corresponds to stocks that are actively traded in both markets. **Figure 1** describes the membership of the grid cells. Trading intensity is defined as follows. *Light* trading describes stocks that trade less than once every hour. *Moderate* trading describes stocks that trade at least once every hour, but less frequently than once every fifteen minutes (fifteen minutes is the finest interval in our tests). *Active* trading describes stocks that trade once every fifteen minutes or more often. We use the median number of daily transactions for 1995 to classify stocks.

In **table 1** we describe the mean and median daily trading volume, number of trades, and trade size for our sample of dually - listed stocks. The numbers in this table are the cross-sectional mean and median of the stock level means. The sample has a mean of 60 transactions per day on the NYSE and 86 transactions per day on the TSE. While these numbers indicate active trading, we find considerable variation across the stocks. For example, Grid 1 stocks trade on average only four times a day on the TSE, and three times a day on the NYSE. Grid 3 stocks record 118 trades per day on the NYSE, but only eight trades per day on the TSE. Conversely, Grid 7 stocks record 83 trades per day on the TSE, but only five trades per day on the NYSE. These results suggest that neither the TSE nor the NYSE absolutely dominates trading for all dually listed stocks. Finally, the 22 Grid 9 stocks trade heavily in both markets, providing a strong *ex ante* setting in which to examine integration. Trade size generally increases as trading activity for the market increases. This is always true for own market trade size (e.g. NYSE trade size increases with NYSE trading). It generally holds for cross-market trade size as well (e.g. TSE trade size increases with NYSE trading).

For the entire sample, intraday trading volume and number of trades are similar in the two markets, and follow the familiar U-shaped pattern on the TSE and the NYSE (see, for instance, Foster and Vishwanathan, 1993). However, we do see a marked difference in relative spreads on the two markets, with spreads on the NYSE being larger than on the TSE. One reason for this discrepancy is that, in 1995, the NYSE tick and the TSE tick were both 1/8 (in the local currency) – consequently, the relative spread on the NYSE was larger. (Results for the separate grids are very similar to the overall results presented here. To conserve space, these results are not presented.)

[10] However, when we use only posted quotes in each interval, we get similar results.

4 Results

In sections 4.1 through 4.4, we provide results for tests on cointegration, the order flow response to deviations from the law of one price, the behavior of the implied exchange rate, and arbitrage opportunities. Instead of providing results for each firm in our sample, we provide results for the median firm in each grid. The median firm is found by ranking all grid members on the basis of aggregate trading volume in the two markets. One concern is that the median firm's results may not be representative of the grid. Therefore, we also present the averages for each of the nine grids.[11]

4.1 Cointegration Tests

In order to show that U.S. dollar prices on the TSE and the NYSE are cointegrated, we first check that these prices are both I(1). The price series are measured at 15–minute intervals through the day, with the most recent quotes being used in each period. We use an augmented Dickey-Fuller (ADF) test with five lags to test for non-stationarity and confirm that both series have a unit root. (For brevity, we do not report these results.) We then check that the residuals from the (long-run) cointegrating regression are stationary. That is, we first estimate the relation between the U.S. dollar prices of stock *j* on the TSE and the NYSE:

$$P_{j,t}^{TSE} S_t = b_{o,j} + b_{1,j} P_{j,t}^{US} + m_{j,t} \qquad \langle 9 \rangle$$

and test for stationarity of the fitted residuals $\hat{m}_{j,t}$ via an ADF test with five lags. (We repeat all of the ADF tests with ten lags; these lead to identical conclusions.)

Table 3 summarizes the results of this preliminary analysis for the median firm in each grid. We find that we cannot reject the hypothesis that $b_1 = 1$ for any stock. While b_1 is above 0.93 in every case, it is closer to unity for the more heavily traded stocks. The intercept, b_0, is significantly above zero in every case, though the point estimate is large only for Grid 1 and Grid 4 stocks. The significance of b_0 necessitates using the residual from the cointegrating regression $\langle 9 \rangle$, $\hat{m}_{j,t}$, in the error correction models, rather than the more appealing price differential. In the ADF test for stationarity of the $\hat{m}_{j,t}$, the coefficient on $\hat{m}_{j,t-1}$ is always less than one. Thus, the residuals from the cointegrating regression are stationary and we can reject the null hypothesis of no cointegration. The difference from one is small for the least actively traded stocks, and it increases with trading activity in either market.

Cointegration implies that we can estimate error correction models for U.S. dollar TSE returns and NYSE returns, as described by $\langle 3a \rangle$ and $\langle 3b \rangle$. **Table 4** provides the fully specified error correction regressions for the median stock from each grid. Table 4 also provides the mean of the error correction coefficients for the nine grids. The coefficient on the error correction term, $\hat{m}_{j,t}$, is positive in the NYSE return regression, and negative in the TSE return regression. This is the case both for the median firm in each grid, and for the corresponding grid average. Thus, when the U.S. dollar price on the TSE exceeds the NYSE price, the TSE return is negative and the NYSE return is positive in the following period. In other words, if prices deviate from the law of one price in any period, there is a reduction in the size of the deviation in subsequent periods, ceteris paribus. This convergence in prices supports the law of one price and market integration.[12]

In the error correction models, the first two lags of own–market returns are almost always significant, though the coefficients are negative.[13] In the model for NYSE returns, the coefficient on the first lag of the TSE return is positive and significant, and higher order lags are also occasionally significant. The

---

[11] The standard error of the mean assumes independence of the coefficients across stocks. However, we do not focus on the statistical significance of the cross-sectional mean.

[12] These results contrast with the weaker results in the PPP literature.

[13] This cannot be due to bid–ask bounce, since we are using midpoint returns. One possible explanation is inventory management by dealers in both markets.

transmission from lagged NYSE to TSE returns is not as pervasive. These coefficient values are important in understanding the patterns in the speed of adjustment that we describe shortly.

The error correction models indicate that prices in the two markets tend to move towards each other in response to shocks that drive them apart, but do not address either the speed of adjustment or the process by which this adjustment takes place. Conclusions about these issues cannot be based solely on the error correction coefficients, since adjustment also depends on the coefficients on the lagged returns in the error correction models. However, we can use the full set of estimated coefficients from the models to study the speed and process of price adjustment.

We study the price adjustment process for the median firm in every grid. We initially set share prices in each market to obey the law of one price. Next, we introduce an artificial price discrepancy, by assigning the residual from ⟨9⟩, $\hat{m}_{j,t}$, a value of two standard deviations (the standard deviation is firm–specific and is intended to represent a large shock). We finally use the parameter estimates from the error correction models (⟨3a⟩ and ⟨3b⟩) to recursively generate returns and new error terms for thirteen successive intervals, which is equivalent to one trading day. (The initial values of the five lagged returns are set to zero.) **Figure 2 and figure 3** summarize the results. Figure 2 displays the errors as a function of elapsed time and figure 3 the corresponding returns in each market. By construction, the shock is +2 standard deviations, that is, TSE prices are temporarily above NYSE prices. Hence, adjustment involves positive returns on the NYSE and negative returns on the TSE.

Two features of the price adjustment process are noteworthy. First, the induced error decays quickly: the half-life is less than thirty minutes in all cases. (The error in interval 1 is the initial deviation from the law of one price.) Price adjustment takes less time as trading activity declines. In grids 1, 2, 3, 4, 5, and 7, the deviation is reduced to close to zero in only two intervals (30 minutes). However, in grids 6, 8, and especially 9, the deviation is different from zero for at least an additional thirty minutes. This pattern can be explained by the inverse relation between liquidity and price changes. With thin markets, orders quickly result in quote revisions.

Second, the decay in the deviation from the law of one price is not monotonic. Every grid is characterized by an oscillatory dampening that suggests over–correction in the two markets. This is seen directly in figure 3, which shows large returns of the predicted sign in interval 2, the period after the shock. However, subsequent returns reverse sign, even though their magnitude shrinks through time. These returns should have been uniformly negative on the TSE and uniformly positive on the NYSE. The reason for the alternating signs can be traced to the signs of the coefficients on the lagged returns in the error correction models, which are positive for cross–market returns and negative for own–market returns. A monotonic pattern in returns would have been assured by exactly the opposite set of coefficient signs. We do not explore the reasons for these coefficient patterns. Despite the oscillatory effects, the deviation from the law of one price decays rapidly for every stock.

Figure 3 shows that as NYSE trading activity increases (we move rightwards in the figure), the TSE generally adjusts to an increasing degree (we see larger returns). Likewise, as TSE trading activity increases, NYSE returns get larger. Figure 3 also shows, however, that more of the adjustment seems to occur on the NYSE compared to the TSE.

4.2 Order flow tests

In order to examine the response of order flow to deviations from the law of one price, we run the VARs described in ⟨4a⟩ and ⟨4b⟩. We use the signed number of trades in each 15-minute interval to measure order flow (using signed volume instead provides similar inferences). We first confirm that order flow is stationary, via ADF tests; the stationarity of order flow allows us to run the VARs. The null hypothesis of a unit root is rejected for every stock (results suppressed).

Table 5 contains the results of the VAR estimation. We focus on the coefficient associated with $\hat{m}_{j,t}$, the deviation from the law of one price in ⟨9⟩. Recall that, in an integrated market, NYSE order flow is expected to be positively associated, and TSE order flow negatively associated, with the error. As predicted, the coefficient on $\hat{m}_{j,t}$ is significantly positive in the NYSE order flow equation in every grid.

However, while the coefficient in the TSE order flow equation is of the predicted sign for eight of the nine stocks, it is insignificantly different from zero for each stock. The mean coefficient in each grid also supports the view that the NYSE order flow response to pricing errors is significant, while the TSE response is not. Examining the coefficient estimates, it is clear that the response of order flow to the error increases with NYSE trading activity. Thus, moving from Grid 1 to Grid 3, or from Grid 4 to Grid 6, or from Grid 7 to Grid 9, the error correction coefficient gets stronger. However, no such pattern is evident as TSE trading activity varies.[14]

Overall, the order flow regressions suggest the NYSE is the venue where error correcting order flow is routed. This conclusion is a bit puzzling: savvy traders should be willing to route some trades to the TSE in the event of price deviations (e.g. sell transactions should migrate to the TSE if the price is higher than the NYSE price). Instead, it appears as if traders are willing to route their trades to the NYSE (but not the TSE) in response to price differences.[15] However, it is unclear why the NYSE dominates price sensitive order flow even for stocks that are lightly traded on the NYSE.[16]

4.3 Implied exchange rate tests

Expression ⟨5⟩ states that the distribution of the implied exchange rate should be centered on, and clustered around, the spot rate. This implied exchange rate is computed by stock from ⟨6⟩ as the ratio of the NYSE price to the TSE price (each measured in local currency).

We measure the percentage deviation between the estimate for each stock and the contemporaneous spot rate. Let $D_{s,i,t}$ be the percentage deviation for stock s in interval i on date t. As a measure of bias, we

[14] The coefficients on lagged order flow in the VARs show the positive influence of the first two lags of own–market order flow on the NYSE for most stocks and on the TSE for each stock. Cross–market order flow is not significant in either equation. After controlling for deviations from the law of one price, therefore, no order flow spillovers exist across markets.

[15] This result, however puzzling, is consistent with anecdotal evidence that institutional traders in Canada often send their trades to the NYSE.

compute the cross-sectional mean of $D_{s,i,t}$ for interval i and date t. (The cross-sectional distribution of $D_{s,i,t}$ should be mean zero.) As a measure of dispersion we compute the root mean square error, RMSE, defined as

$$s_{i,t} = \sqrt{\sum_{s=1}^{N} \Delta^2_{s,i,t} / N} \qquad \langle 10 \rangle$$

We calculate the RMSE instead of a conventional standard deviation since we are interested in the dispersion around the *true* spot rate, and wish to avoid incorporating the mean bias in the calculation. (In fact, using the conventional standard deviation reduces the dispersion by 10%.)[17] **Table 6** contains the time series mean of the cross-sectional mean and the RMSE of the distribution of $D_{s,i,t}$. (Normalization by the spot rate is intended to make the time series observations homoscedastic.) We present these statistics by hourly interval since, paralleling the time of day effects documented in table 2, there may be systematic differences by time of day.

Panel A of table 6 and **figure 4** present the mean bias. The results indicate some bias for the most actively traded (Grid 9) stocks, which varies between –1 and –7 basis points. The bias is largest at the end of the day, almost 7 basis points. However, this bias is not large and its significance is driven by the tiny standard error rather than by a large point estimate. In fact, the point estimate of the bias for Grid 9 stocks is smaller than the bias for the other stocks. In economic terms, the bias is small, less than the bid-ask spread in the inter-bank market.

Panel B of table 6 and **figure 5** show that the mean RMSE is less than 50 basis points for Grid 9 stocks. The dispersion increases more or less steadily as we move to less actively traded stocks, reaching a maximum of 200 basis points for Grid 1 stocks. However, the dispersion is within reasonable transaction cost bounds for most stocks. Using a two standard deviation bound to identify extreme

[16] We also estimated four–variable VARs, modeling returns and order flow in the two markets. Our conclusions regarding price adjustment and order flow are unaltered.

[17] For Grid 3, where we have only one stock, the RMSE reduces to the absolute value of $D_{s,i,t}$.

observations, a deviation greater than one percent from the spot rate is unusual for Grid 9 stocks, but a deviation as large as four percent for Grid 1 stocks is not. Given reasonable estimates of transactions costs, even extreme deviations are unlikely to allow arbitrage. For most stocks, a relative bid-ask spread of one percent in each market will prevent arbitrage. Of course, traders seeking the lowest cost market to transact may be able to reap savings as large as four percent, excluding transactions costs (if they are trading in Grid 1 stocks).

There is a small tendency for the RMSE to be higher at the open and close. The intraday U-shape to the cross-sectional RMSE can perhaps be explained by price pressure effects in the market for individual equities. Specifically, heightened trading at the beginning and end of the day shows up in relatively large deviations from the true spot rate for a few stocks. Even at these times, the deviations for most of the stocks in our sample are small and consistent with integration.

We check that our results are not mechanically produced by two effects. First, if quotes for some of the stocks are stale, the RMSE will be inflated. When we re–estimate the RMSE using quotes for stocks that have been revised in that interval, the RMSE is very similar. Second, since there are fewer stocks in the lower grids, the RMSE could be less precise, and the time series distribution could contain more extreme values. To confirm that this effect is not behind the larger RMSE for the lower grids, we aggregate the grids. A natural aggregation is grids 1, 2, 4, 5 (24 stocks); grids 3, 6, 7 and 8 (22 stocks); and grid 9 (22 stocks).[18] If differing numbers of stocks drive the greater dispersion for the less actively traded stocks, the RMSE should be similar for the three aggregates. However, the RMSE increases with trading activity for the coarser grids.

4.4 Arbitrage opportunities

Figure 6 and figure 7 summarize the time–series distributions of percentage deviations between the

[18] The three groups correspond to moderate activity or less in each market, active trading in one but not the other, and active trading in both, respectively.

implied and actual spot rates for the median stock in each grid.[19] Figure 6 is based on 〈8〉, and directly measures the percentage returns to an arbitrage strategy, excluding spreads and other costs. Figure 6 shows that the magnitude of the profits declines as trading activity in either market increases, for example as we move from Grid 1 to Grid 5 and then on to Grid 9. For actively traded stocks, the variation in the rates of return is small, of the order of one percent for Grid 9, and with few outliers.

In Figure 7, we directly measure the relative frequency of arbitrage opportunities for the median stock in each grid. We calculate the incidence of arbitrage opportunities using 〈7a〉 and 〈7b〉 over the entire sample period, explicitly taking into account the bid and ask spreads in the stock and the foreign exchange markets. Overall, our sample of dually listed stocks presents limited arbitrage opportunities. The most arbitrage prone stock is in grid 4: for this stock, approximately three (two) of every 100 intervals presented arbitrageable quotes (net of 2% commissions). For all other grids, the incidence of arbitrage opportunities net of 2% commissions is approximately one percent or less. We interpret these results as follows. Grid 1 stocks offer few arbitrage opportunities chiefly because of the higher bid-ask spreads in these stocks. Grid 9 stocks have low spreads, but offer few arbitrage opportunities because of active trading in both markets, which keeps prices in a tight bound around the spot rate.

5 Concluding comments

In this paper, we use intraday data for a sample of 68 Canadian stocks that trade on both the Toronto Stock Exchange (TSE) and the New York Stock Exchange (NYSE) to examine the extent of integration of the TSE and the NYSE. Our experiment is distinctive for several reasons: these stocks trade as stocks (not ADRs) in the U.S.; the incremental transactions costs of trading abroad are not large, especially for institutional traders; and the markets have completely synchronous trading. Additionally, we sort the

[19] Alberta Energy Corp. only started trading on the NYSE in September. As a result, there are only three months of data for this stock. We decided to stay with this firm since it was the original selection for Grid 7.

sample of stocks according to the level of trading activity on the TSE and the NYSE. The markets for the heavily traded stocks are expected to be integrated.

An implication of market integration is that prices on the TSE and the NYSE will differ by the Canadian dollar–U.S. dollar exchange rate (the law of one price). We conduct several tests of the law of one price. Our results show that prices in the two markets are cointegrated. Error correction model estimates show rapid adjustments in both markets to deviations from the law of one price. In general, the extent of a market's price adjustment declines as trading activity in that market increases. The ratio of the NYSE price to the TSE price, the *implied* spot exchange rate, is tightly distributed in cross–section around the actual spot rate, allowing few opportunities for arbitrage. The time–series of returns to arbitrage positions does not indicate substantial or frequent profit opportunities. While the results are strongest for actively traded stocks, the same pattern is evident for all stocks. Overall, these results point to markets where prices are well integrated.

Surprisingly, the order flow response to price deviations occurs solely on the NYSE. In other words, price sensitive traders in both markets appear to avoid the TSE and route trades to the NYSE. Given, for instance, a case where the U.S. dollar TSE price is greater than the NYSE price, we would expect sell orders from investors in both countries to flow to the TSE. The fact that these transactions do not appear to reach the TSE, even for actively traded stocks, is puzzling, and is at odds with the rest of our evidence, which favors integrated trading.

References

Dornbusch, R. (1987), Purchasing Power Parity, *The New Palmgrave: A Dictionary of Economics*, 1075-1085.

Engel, C. (1993), Real Exchange Rates and Relative Prices? An Empirical Investigation, *Journal of Monetary Economics* 32, 35-50.

Engle, R. and C. Granger (1987), Co-Integration and Error Correction: Representation, Estimation, and Testing, *Econometrica* 55, 251-276.

Foster, D. and S. Vishwanathan (1993), Variations in Trading Volume, Return Volatility, and Trading Costs: Evidence on Recent Price Formation Models, *Journal of Finance* 48, 187-211.

Froot, K., M. Kim and K. Rogoff (1995), The Law of One Price over 700 Years, *NBER working paper*.

Harris, F. H. deB, T. McInish, G. Shoesmith and R. Wood (1995), Cointegration, Error Correction, and Price Discovery on Informationally Linked Security Markets, *Journal of Financial and Quantitative Analysis* 30, 563-579.

Harvey, Campbell (1991), The World Price of Covariance Risk, *Journal of Finance* 46, 111-157.

Hasbrouck, J. (1995), One Security, Many Markets: Determining the Contributions to Price Discovery, *Journal of Finance* 50, 1175-1200.

Ip, G. (1998), Proposed Nasdaq-Amex Deal May Prompt Other Mergers, *The Wall Street Journal, Money and Investing Section* (March 16, 1998).

Ito, T., R. Lyons, and M. Melvin (1998), Is There Private Information in the FX Market? The Tokyo Experiment, *Journal of Finance* 53, 1111-1130.

Jorion, P. and E. Schwartz (1986), Integration v. segmentation in the Canadian stock market, *Journal of Finance* 41, 603-614.

Karolyi, A. (1995), A Multivariate GARCH Model of International Transmissions of Stock Returns and Volatility: The Case of the United States and Canada, *Journal of Business and Economic Statistics* 13, 11-25.

Kleidon, A. and I. Werner (1996), U.K. and U.S. trading of British cross-listed stocks: An intraday analysis of market integration, *Review of Financial Studies* 9, 619-664.

Lee, C. and M. Ready (1991), Inferring Trade Direction from Intraday Data, *Journal of Finance* 46, 733-746.

Lieberman, O., U. Ben-Zion and S. Hauser (1999), A Characterization of the Price Behavior of International Dual Stocks: An Error Correction Approach, *Journal of International Money and Finance* 18, 289-304.

Neumark, D., P. Tinsley and S. Tosini (1991), After-Hours Stock Prices and Post-Crash Hangovers, *Journal of Finance* 46, 159-178.

Rogoff, K. (1996), The Purchasing Power Parity Puzzle, *Journal of Economic Literature* 34, 647-668.

Table 1

Daily Volume, Number of Trades, and Trade Size for 68 TSE stocks cross listed on the NYSE/AMEX. The sample period is Jan 3, 1995 to Dec 29, 1995.

The top row in each cell is the grand mean; the row below is median of the stock-level means.

	Volume		No. Trades		Trade Size	
	TSE	NYSE	TSE	NYSE	TSE	NYSE
All Stocks	218,012	172,936	86	60	2,894	2,449
N=68	148,650	132,358	75	53	1,466	1,687
Grid 1	12,418	4,481	4	3	3,602	1,139
N=4	1,622	1,550	3	2	531	641
Grid 2	11,268	29,551	5	14	1,899	1,967
N=8	2,850	16,475	4	12	629	1,301
Grid 3	21,812	244,522	8	118	2,579	2,029
N=1	7,305	195,450	7	112	966	1,734
Grid 4	66,881	7,808	16	3	4,547	2,258
N=9	15,750	2,478	14	3	1,139	848
Grid 5	79,436	46,722	22	19	3,660	2,316
N=3	32,768	29,433	19	15	1,524	1,583
Grid 6	58,505	188,204	17	52	3,060	3,454
N=4	22,436	123,750	13	44	1,503	2,792
Grid 7	207,337	10,386	83	5	2,459	1,446
N=8	136,800	3,188	73	4	1,528	741
Grid 8	356,909	62,605	125	20	3,023	2,734
N=9	246,275	36,094	108	17	1,950	1,749
Grid 9	396,276	438,683	170	152	2,436	3,043
N=22	292,315	350,498	148	135	1,861	2,491

Table 2

Average Daily Volume, Number of Trades, and Relative Bid-Ask Spread by time of day. Interval 1 corresponds to 9:30-10:00 a.m., while the other intervals are the next six hours in the trading day (10:00-11:00 a.m., 11:00 a.m.-12:00 p.m.,.., 3:30-4:00 p.m.). Volume and number of trades for interval 1 have been multiplied by two.

The sample consists of 68 TSE stocks cross listed on the NYSE/AMEX. The sample period is Jan 3, 1995 to Dec 29, 1995.

	Volume		No. of Trades		Rel Spread	
Interval	TSE	NYSE	TSE	NYSE	TSE	NYSE
1	65,354	45,242	26.5	16.2	0.02127	0.04950
2	39,854	27,971	14.4	10.1	0.01946	0.04026
3	32,822	24,156	12.5	8.6	0.01883	0.03833
4	23,901	19,014	10.0	7.0	0.01865	0.03782
5	19,369	16,828	9.1	6.7	0.01845	0.03725
6	24,886	20,045	11.1	7.7	0.01838	0.03742
7	38,487	24,762	14.9	9.4	0.01814	0.03722

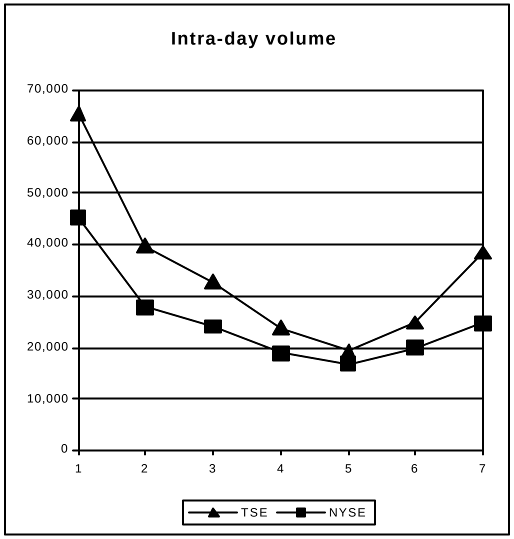

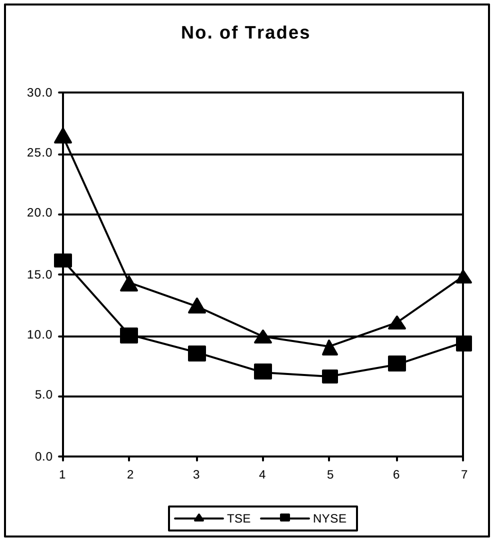

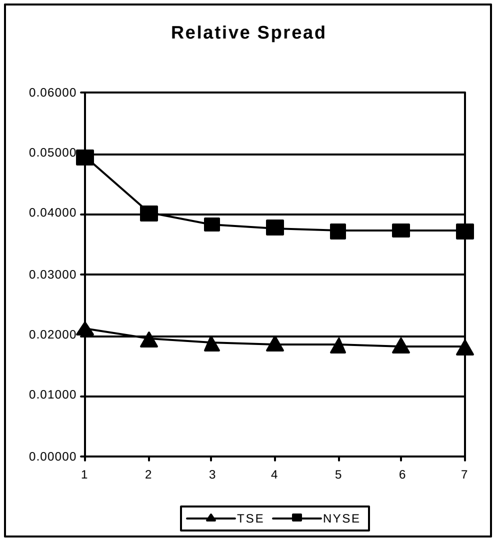

Table 3

B_0 and B_1 are the intercept and the slope coefficient in the regression of the US$ TSE price on the contemporaneous NYSE price for the median firm in each grid. *Nu* is the coefficient on the lag 1 error in the ADF test, and should be significantly below one for stationarity. Standard errors are reported below the coefficient in italics. The median firm is selected from a volume-ranked list of all firms in the grid.

Grid	Ticker	B_0	B_1	Nu	Grid Average B_1	N
1	CMW	0.1620	0.9306	-0.0416	0.9720	3
		0.0071	*0.0032*	*0.0045*	*0.0226*	
2	CID	0.0025	0.9990	-0.1114	0.9922	7
		0.0021	*0.0008*	*0.0093*	*0.0045*	
3	HSM	0.0254	0.9903	-0.2080	0.9903	1
		0.0036	*0.0014*	*0.0139*		
4	RJL	0.1308	0.9431	-0.1001	0.9884	8
		0.0059	*0.0026*	*0.0077*	*0.0116*	
5	UDI	0.0135	0.9955	-0.1523	0.9368	2
		0.0015	*0.0005*	*0.0106*	*0.0398*	
6	PGU	0.0186	0.9927	-0.3485	0.9535	3
		0.0021	*0.0008*	*0.0194*	*0.0442*	
7	AEC	0.0167	0.9936	-0.1528	0.9986	8
		0.0100	*0.0036*	*0.0190*	*0.0035*	
8	DTC	0.0050	0.9976	-0.3029	0.9970	9
		0.0016	*0.0007*	*0.0169*	*0.0014*	
9	BCE	0.0169	0.9952	-0.4984	0.9970	21
		0.0028	*0.0008*	*0.0250*	*0.0008*	

Table 4

Coefficients from the Error Correction Model (ECM) are presented for the median (volume-ranked) stock in each grid (TSE ticker in parentheses). *Error Term* refers to the residual from the cointegrating regression. In addition, the last column presents the average of the coefficients for the grid. The ECM is based on 15-minute returns. $R_{TSE,t}$ ($R_{US,t}$) refers to TSE (NYSE) returns. Incr R^2 refers to the incremental R^2 associated with the addition of the error term to the ECM. Coefficients significant at the five percent level are in bold.

Panel A: Dependent variable = R_{TSEt}

GRID	INTER-CEPT	ERROR TERM	$R_{TSE,t-1}$	$R_{TSE,t-2}$	$R_{TSE,t-3}$	$R_{TSE,t-4}$	$R_{TSE,t-5}$	$R_{US,t-1}$	$R_{US,t-2}$	$R_{US,t-3}$	$R_{US,t-4}$	$R_{US,t-5}$	Adj-R^2	Incr.R^2	Grid mean Error term
1 (CMW)	0.0000	**-0.013**	**-0.079**	**-0.041**	-0.016	**-0.027**	**0.031**	**0.042**	**0.070**	**-0.043**	0.011	-0.001	0.016	0.002	-0.019
2 (CID)	0.0000	**-0.092**	**-0.249**	**-0.111**	**-0.061**	**-0.056**	-0.018	**0.136**	**0.052**	0.012	0.025	0.031	0.111	0.020	**-0.076**
3 (HSM)	0.0000	**-0.130**	**-0.053**	-0.016	**-0.037**	0.008	0.023	0.010	**0.026**	**0.040**	0.012	-0.009	0.065	0.030	-0.130
4 (RJL)	0.0000	**-0.004**	**-0.064**	**-0.040**	0.000	**-0.029**	**-0.039**	0.009	0.013	-0.004	-0.013	**-0.020**	0.009	0.000	**-0.020**
5 (UDI)	0.0000	**-0.039**	**-0.038**	0.000	-0.001	0.027	-0.003	0.021	0.016	0.012	0.021	0.014	0.009	0.003	-0.036
6 (PGU)	0.0000	**-0.135**	0.016	-0.003	-0.011	-0.008	0.010	-0.020	-0.017	0.006	-0.006	0.011	0.034	0.017	-0.097
7 (AEC)	0.0000	**-0.017**	**-0.122**	-0.023	0.009	-0.025	-0.029	0.059	0.033	0.033	0.011	0.021	0.012	0.000	-0.006
8 (DTC)	0.0000	**-0.067**	**-0.066**	**-0.036**	-0.004	0.010	0.018	**0.055**	0.013	0.013	0.002	0.006	0.024	0.005	**-0.055**
9 (BCE)	0.0000	**-0.068**	-0.042	0.020	-0.033	-0.003	-0.025	0.010	0.014	0.031	-0.005	0.016	0.010	0.002	**-0.119**

Panel B: Dependent variable = $R_{US,t}$															
GRID	INTER-CEPT	ERROR TERM	$R_{TSE,t-1}$	$R_{TSE,t-2}$	$R_{TSE,t-3}$	$R_{TSE,t-4}$	$R_{TSE,t-5}$	$R_{US,t-1}$	$R_{US,t-2}$	$R_{US,t-3}$	$R_{US,t-4}$	$R_{US,t-5}$	Adj-R^2	Incr.R^2	Grid mean Error term
1 (CMW)	0.0000	0.028	0.131	0.037	0.010	0.034	0.007	-0.036	0.006	-0.028	-0.017	-0.005	0.042	0.011	0.033
2 (CID)	0.0000	0.036	0.020	0.019	0.042	0.004	0.008	-0.038	-0.015	-0.043	-0.002	0.012	0.015	0.006	0.069
3 (HSM)	0.0000	0.081	0.152	0.079	0.010	0.051	0.039	-0.305	-0.142	-0.088	-0.047	-0.040	0.112	0.005	0.081
4 (RJL)	0.0000	0.103	0.072	0.074	0.051	0.005	0.032	-0.075	-0.057	-0.034	-0.049	-0.010	0.069	0.033	0.067
5 (UDI)	0.0000	0.129	0.089	0.071	0.053	0.042	0.032	-0.108	-0.108	-0.060	-0.017	-0.018	0.089	0.029	0.099
6 (PGU)	0.0000	0.222	0.196	0.134	0.069	0.088	0.023	-0.270	-0.170	-0.094	-0.091	-0.019	0.158	0.018	0.108
7 (AEC)	0.0000	0.139	0.004	0.019	0.052	0.040	-0.010	-0.057	-0.046	0.011	0.000	-0.040	0.090	0.045	0.138
8 (DTC)	0.0000	0.254	0.132	0.023	0.072	0.054	0.017	-0.125	-0.062	-0.049	-0.034	0.013	0.148	0.042	0.210
9 (BCE)	0.0000	0.468	0.131	0.095	0.033	0.042	0.011	-0.161	-0.086	-0.028	-0.036	-0.008	0.218	0.044	0.338

Table 5

The response of order flow (signed number of trades) to deviations from the law of one price. Coefficients from the regression are presented for the median (volume-ranked) stock in each grid (TSE ticker in parentheses). *Error Term* refers to the residual from the cointegrating regression. In addition, the last column presents the average of the coefficients for the grid. The regression is carried out over 15-minute intervals. $OF_{TSE,t}$ ($OF_{US,t}$) refers to TSE (NYSE) order flow. Incr R^2 refers to the incremental R^2 associated with the addition of the error term to the regression.

Coefficients significant at the five percent level are in bold.

Panel A: Dependent variable = $OF_{TSE,t}$

GRID	INTER-CEPT	ERROR TERM	$OF_{TSE,t-1}$	$OF_{TSE,t-2}$	$OF_{TSE,t-3}$	$OF_{TSE,t-4}$	$OF_{TSE,t-5}$	$OF_{US,t-1}$	$OF_{US,t-2}$	$OF_{US,t-3}$	$OF_{US,t-4}$	$OF_{US,t-5}$	Adj-R^2	Incr.R^2	Grid mean Error term
1 (CMW)	**-0.091**	-0.530	**0.049**	**0.025**	0.008	**0.030**	0.011	0.027	-0.016	0.028	0.030	0.026	0.004	0.000	-0.382
2 (CID)	**-0.157**	1.444	**0.047**	**0.025**	0.013	**0.031**	0.015	**-0.046**	0.002	-0.016	-0.026	-0.018	0.006	0.000	0.015
3 (HSM)	**-0.252**	-1.084	**0.114**	**0.026**	**0.043**	0.018	0.011	-0.005	-0.005	0.000	0.002	-0.003	0.018	-0.001	-1.084
4 (RJL)	**-0.446**	-0.561	**0.101**	**0.037**	-0.001	0.021	0.009	0.098	-0.072	0.098	0.031	-0.049	0.013	-0.001	0.380
5 (UDI)	**-0.518**	-8.028	**0.107**	**0.043**	**0.024**	**0.032**	**0.047**	-0.056	**-0.068**	-0.030	-0.018	0.020	0.023	-0.001	-4.802
6 (PGU)	**-0.435**	-0.441	**0.131**	**0.089**	**0.050**	-0.006	**0.025**	**-0.018**	0.003	0.000	-0.007	-0.011	0.036	-0.001	-0.752
7 (AEC)	**-1.798**	-4.638	**0.135**	**0.097**	**0.057**	0.039	**0.080**	-0.278	0.055	0.289	0.168	**0.457**	0.056	0.008	-16.227
8 (DTC)	**-1.487**	-4.194	**0.162**	**0.090**	**0.084**	**0.041**	0.014	-0.009	**0.109**	0.028	-0.079	-0.056	0.062	-0.006	-10.852
9 (BCE)	**-10.311**	-71.315	**0.253**	**0.096**	**0.073**	**0.063**	**0.039**	0.022	0.022	0.033	-0.126	-0.163	0.129	-0.017	-16.730

Panel B: Dependent variable = $OF_{US,t}$															
GRID	INTER-CEPT	ERROR TERM	$OF_{TSE,t-1}$	$OF_{TSE,t-2}$	$OF_{TSE,t-3}$	$OF_{TSE,t-4}$	$OF_{TSE,t-5}$	$OF_{US,t-1}$	$OF_{US,t-2}$	$OF_{US,t-3}$	$OF_{US,t-4}$	$OF_{US,t-5}$	Adj-R^2	Incr.R^2	Grid mean Error term
1 (CMW)	**-0.007**	**0.837**	**0.017**	-0.005	-0.001	-0.009	0.003	**0.026**	0.023	0.003	**0.029**	0.002	0.004	0.003	0.451
2 (CID)	**0.034**	**4.008**	0.002	0.016	0.012	0.003	0.011	**0.072**	**0.025**	0.023	**0.024**	**0.033**	0.012	0.003	**4.979**
3 (HSM)	**-0.254**	**55.014**	0.023	**-0.203**	0.007	0.018	**-0.112**	**0.184**	**0.092**	**0.074**	0.022	**0.043**	0.082	0.010	55.014
4 (RJL)	**-0.021**	**1.900**	**-0.006**	-0.002	-0.001	-0.002	-0.001	0.020	0.004	-0.007	0.007	-0.020	0.012	0.011	**1.544**
5 (UDI)	**0.023**	**7.059**	-0.002	-0.008	-0.007	-0.004	-0.007	**0.096**	**0.066**	**0.047**	**0.051**	**0.028**	0.027	0.003	5.881
6 (PGU)	**-0.222**	**37.610**	**-0.062**	**-0.056**	-0.021	-0.027	-0.003	**0.138**	**0.042**	**0.025**	0.021	**0.027**	0.039	0.010	13.790
7 (AEC)	**0.043**	**13.643**	0.003	0.000	0.003	0.002	0.003	0.011	-0.021	0.039	**0.074**	0.014	0.035	0.027	**8.598**
8 (DTC)	0.031	**19.515**	0.002	0.001	**-0.007**	-0.006	0.002	**0.067**	**0.069**	0.007	**0.033**	**0.036**	0.025	0.014	**16.259**
9 (BCE)	**-0.139**	**147.492**	0.001	0.000	**-0.004**	-0.003	-0.001	**0.113**	**0.056**	**0.053**	**0.056**	**0.067**	0.056	0.017	**90.130**

Table 6

Intra-day bias and dispersion in the implied exchange rate. The implied exchange rate is computed for each stock at the end of every hour as the ratio of contemporaneous NYSE and TSE prices. Bias and dispersion are measured relative to the C$-US$ spot exchange quote. The table presents the mean bias and standard deviation for each interval across the sample period (Jan 3, 1995 to Dec 29, 1995). The ratio for each interval equals zero under complete price integration. Coefficients significant at the 5% level are in bold. Rows correspond to intervals and columns to grids. Interval 1 corresponds to 9:30-10:00 a.m., while the other intervals are the next six hours in the trading day (10:00-11:00 a.m., 11:00 a.m.-12:00 p.m.,.., 3:30-4:00 p.m.).

Interval	Grid								
	9	8	7	6	5	4	3	2	1
PANEL A: BIAS									
1	0.00007	0.00033	0.00012	0.00028	0.00058	-0.00053	-0.00005	0.00186	-0.00341
2	-0.00019	0.00006	0.00084	0.00013	0.00370	-0.00163	-0.00017	0.00018	0.00097
3	**-0.00021**	-0.00012	-0.00051	-0.00029	0.00266	-0.00032	-0.00054	0.00054	-0.00156
4	**-0.00027**	-0.00031	0.00055	-0.00078	0.00438	-0.00041	-0.00060	0.00110	-0.00700
5	**-0.00029**	-0.00026	0.00042	0.00060	0.00259	-0.00263	0.00013	0.00135	-0.00099
6	**-0.00034**	-0.00019	0.00089	0.00076	0.00095	-0.00254	-0.00026	0.00132	-0.00010
7	**-0.00065**	-0.00104	-0.00016	-0.00032	-0.00005	0.00140	-0.00020	-0.00019	-0.00687
PANEL B: DISPERSION									
1	0.00474	0.00728	0.00558	0.00743	0.00782	0.01183	0.00415	0.01460	0.01604
2	0.00445	0.00695	0.00498	0.00676	0.00804	0.01099	0.00385	0.01167	0.01687
3	0.00419	0.00708	0.00477	0.00702	0.00822	0.01044	0.00405	0.01134	0.01356
4	0.00417	0.00699	0.00507	0.00639	0.00548	0.01074	0.00396	0.01073	0.01224
5	0.00420	0.00685	0.00467	0.00716	0.00597	0.01053	0.00385	0.01083	0.01383
6	0.00426	0.00696	0.00451	0.00724	0.00633	0.01061	0.00379	0.01148	0.01778
7	0.00463	0.00747	0.00462	0.00674	0.00575	0.01164	0.00455	0.01198	0.02006

FIGURE 1

Distribution by number of trades of 68 TSE stocks cross listed on the NYSE/AMEX. The sample period is Jan 3, 1995 to Dec 29, 1995.
LIGHT trading is defined as less than one trade every hour on the local exchange; ACTIVE trading is at least one trade every 15 minutes; and MODERATE trading is at least one trade every hour, but less than one trade every 15 minutes. The median number of daily transactions for 1995 is used to classify stocks. *N* refers to the number of firms in each category and is provided below the Grid identifiers.

TRADING ACTIVITY		NYSE / AMEX		
		LIGHT	MODERATE	ACTIVE
TSE	LIGHT	Grid 1 N=4	Grid 2 N=8	Grid 3 N=1
	MODERATE	Grid 4 N=9	Grid 5 N=3	Grid 6 N=4
	ACTIVE	Grid 7 N=8	Grid 8 N=9	Grid 9 N=22

Figure 2

DECAY OF A TWO STANDARD DEVIATION ERROR FROM THE LAW OF ONE PRICE.

TIME IS MEASURED IN 15 MINUTE INTERVALS.

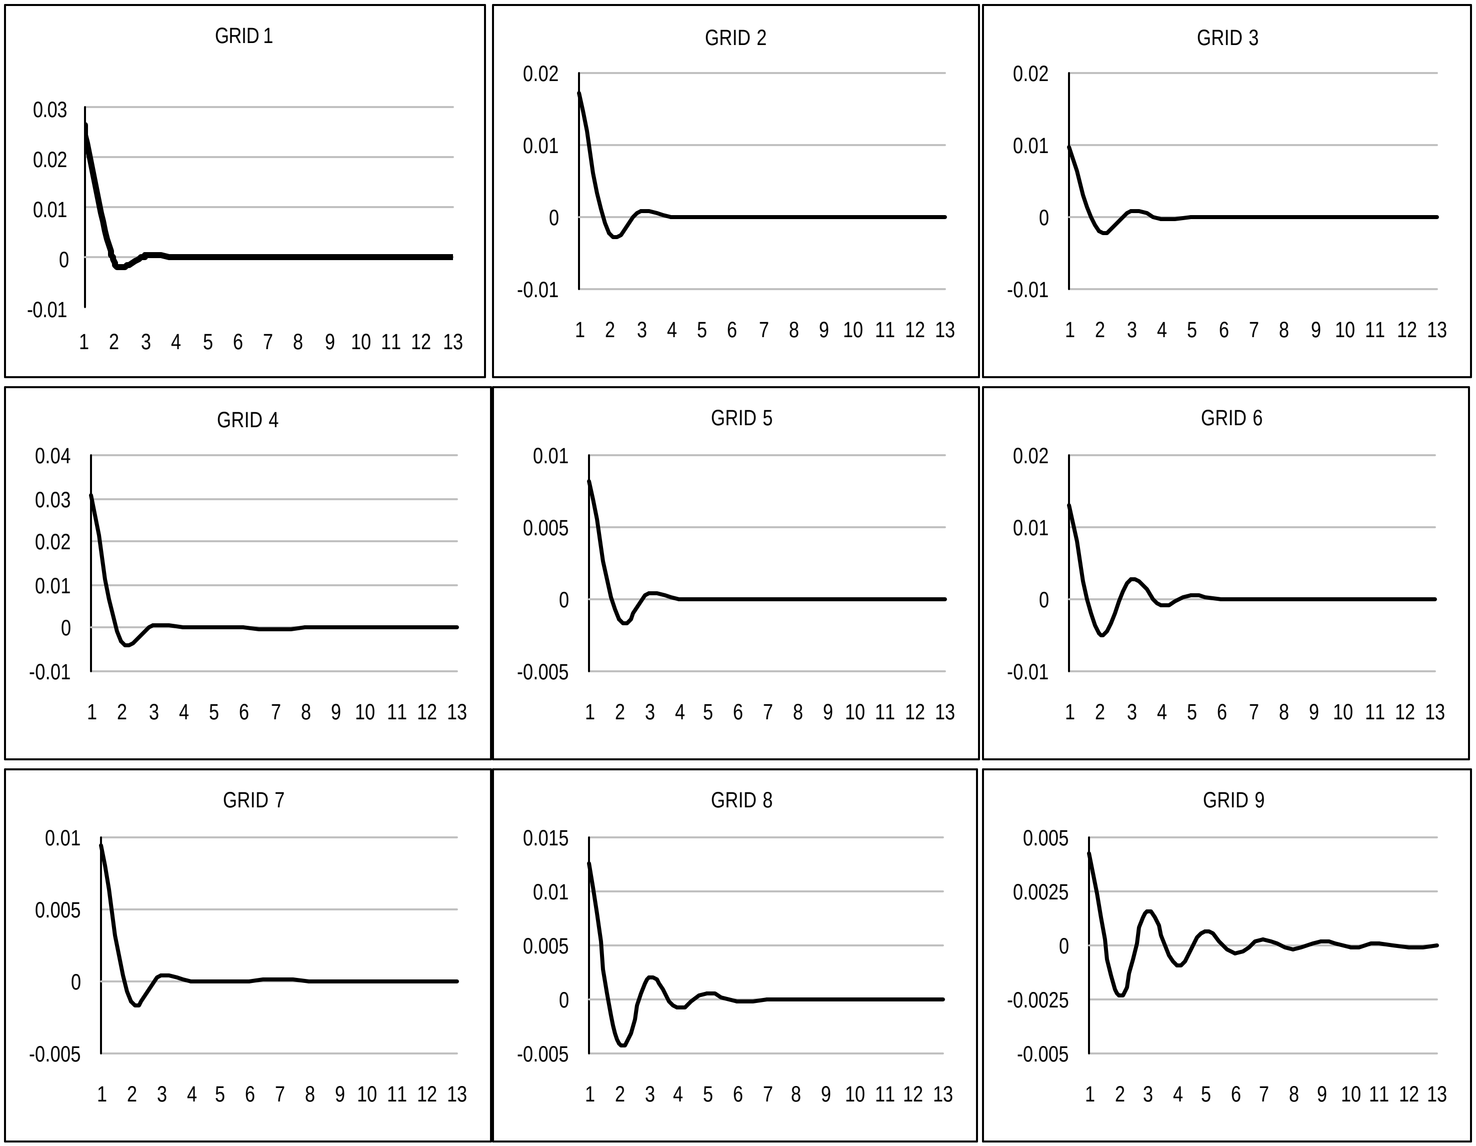

Figure 3

RETURNS ON THE TSE (BLACK) AND THE NYSE (GREY) IN RESPONSE TO DEVIATIONS FROM THE LAW OF ONE PRICE

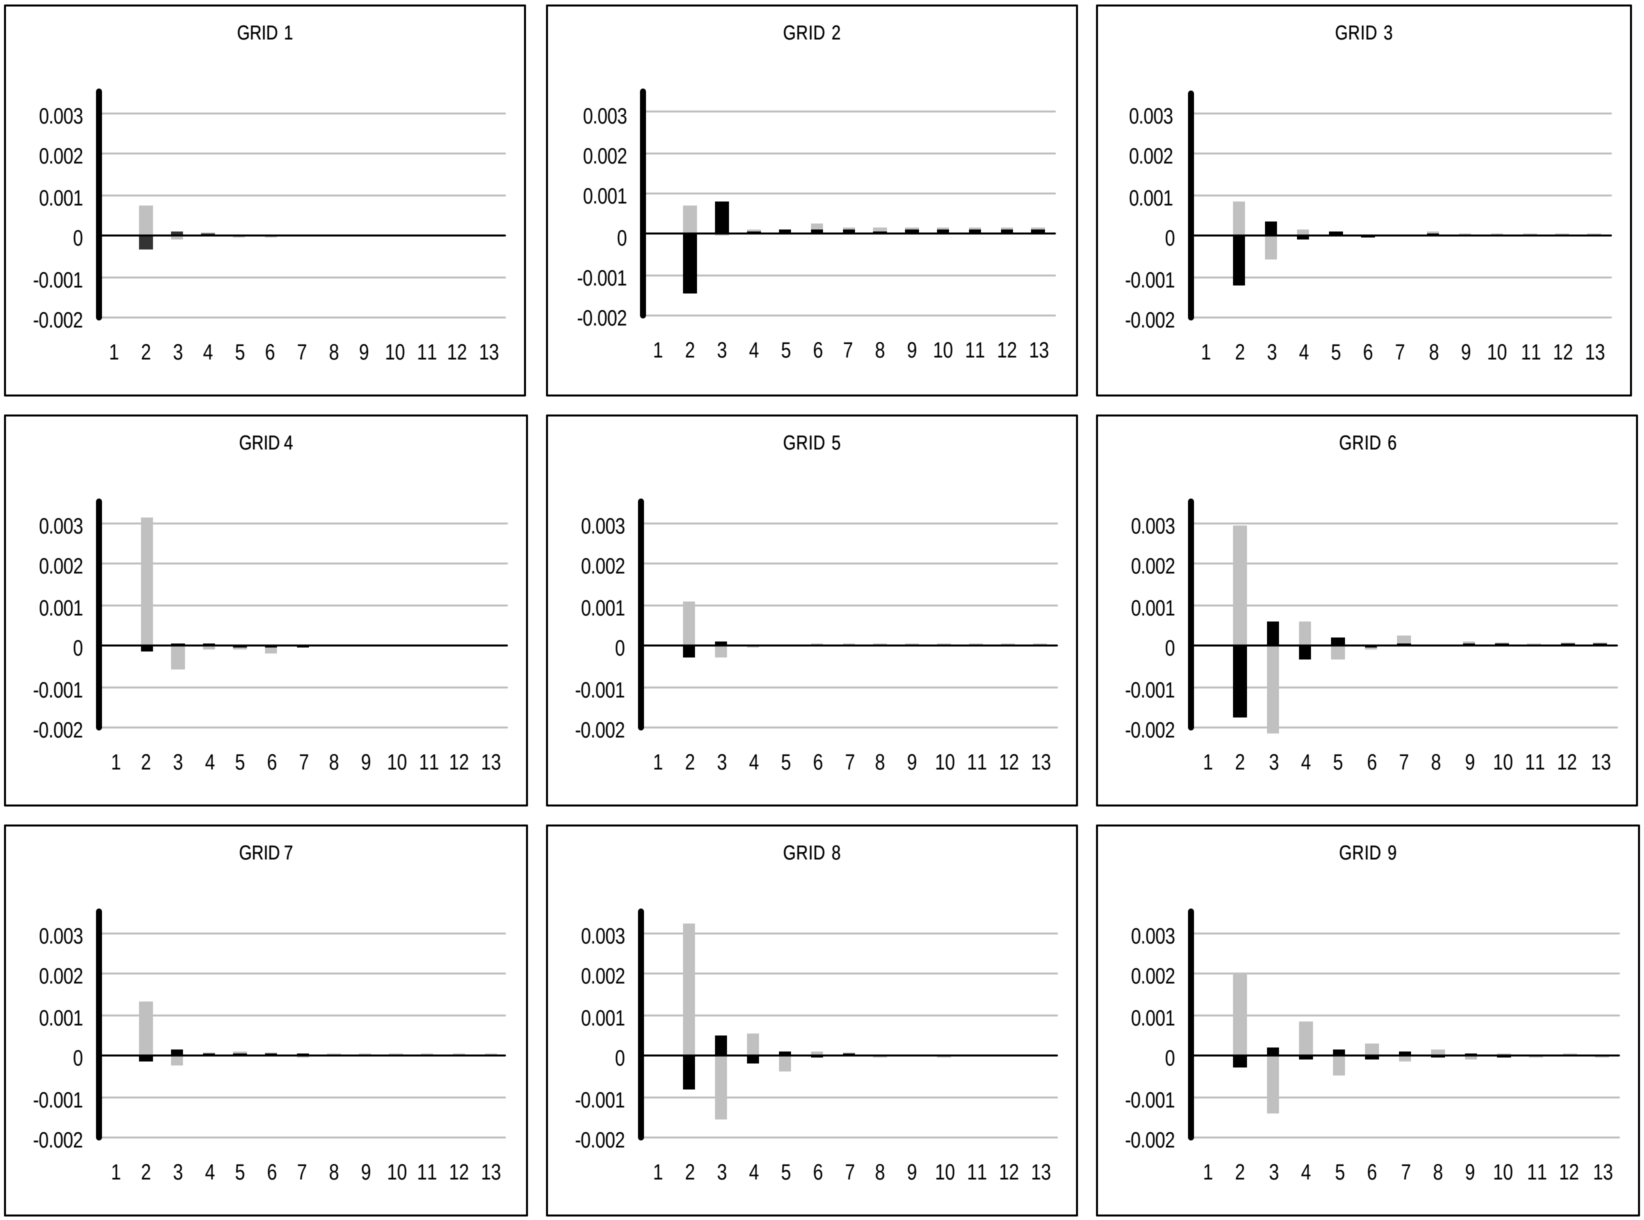

FIGURE 4

Bias in the implied exchange rate

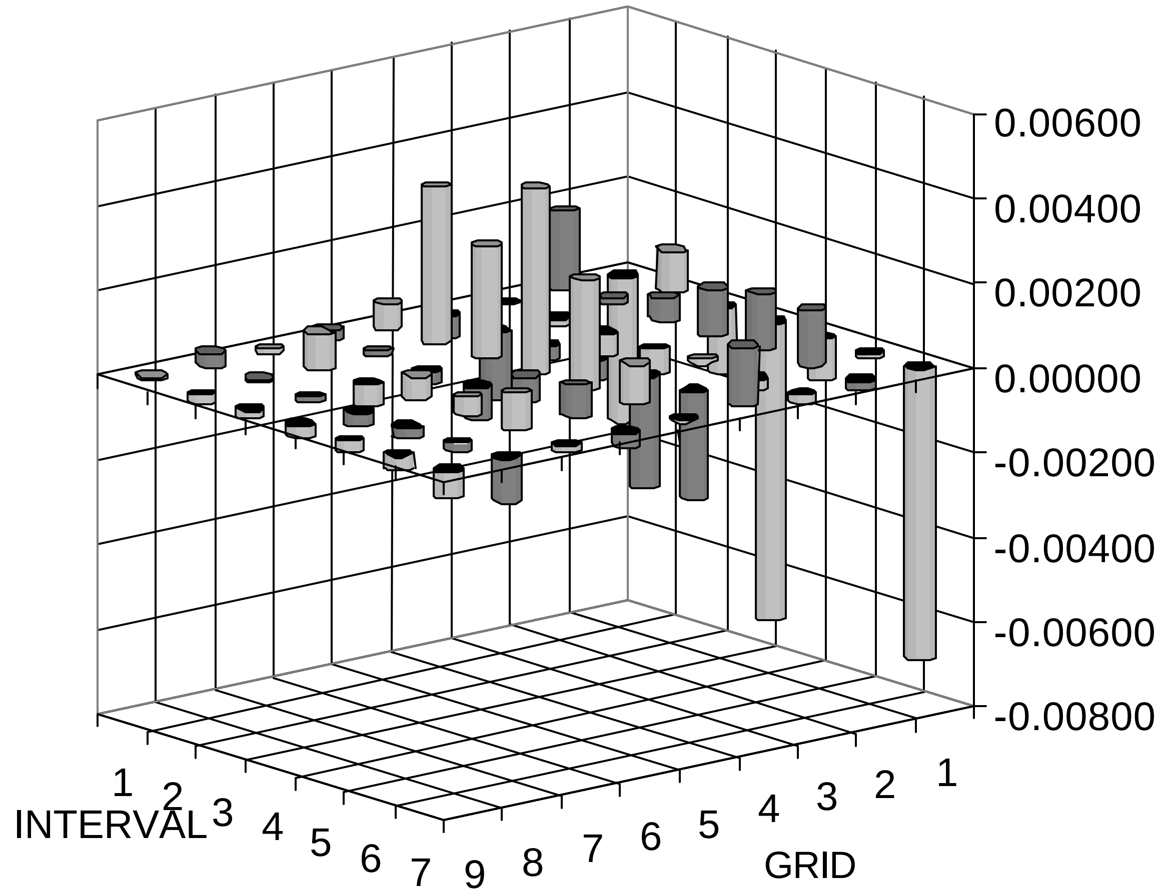

FIGURE 5
Dispersion in the implied exchange rate

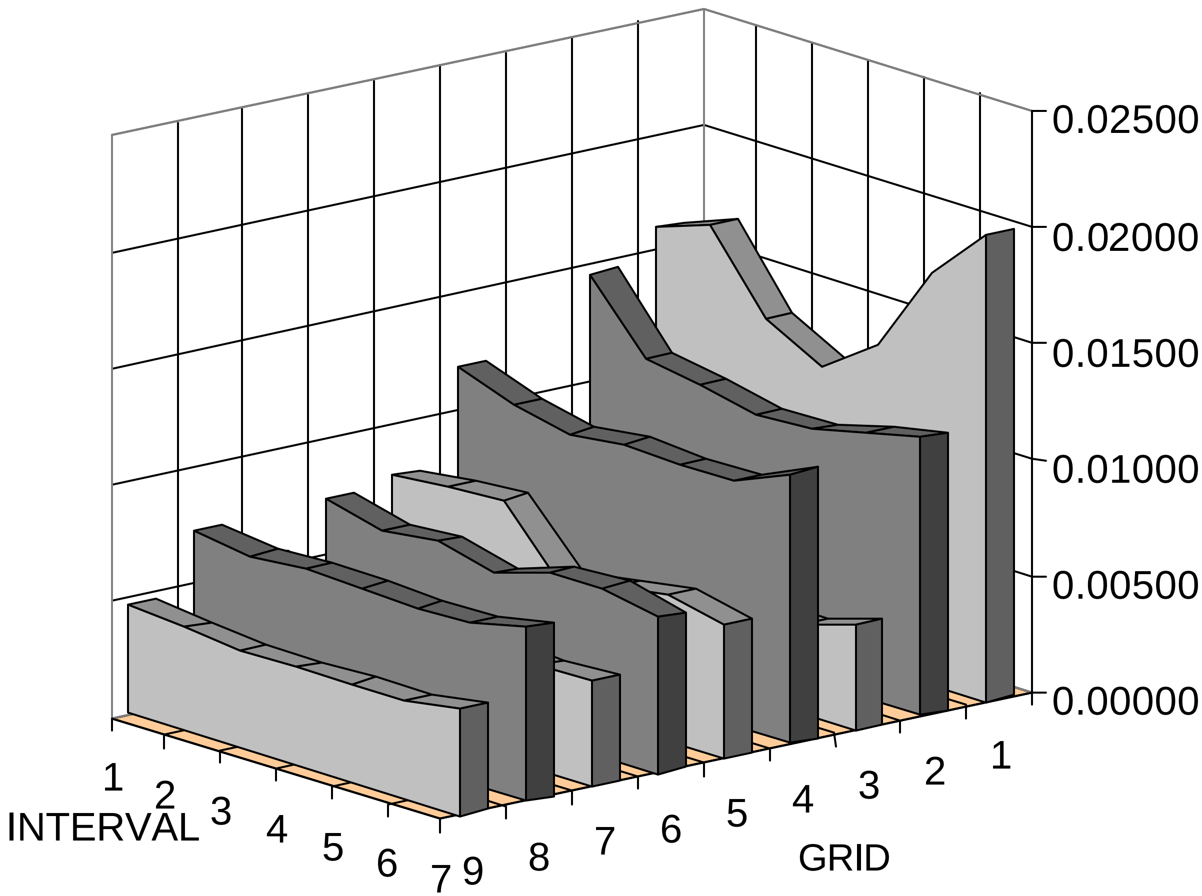

Figure 6

The time series of proportional deviations between the ratio of NYSE to TSE prices and the USD/$ spot exchange rate for the median stock in each grid. The sample period is Jan 3, 1995 to Dec 29, 1995. By construction, the ratio for each interval equals zero under complete price integration.

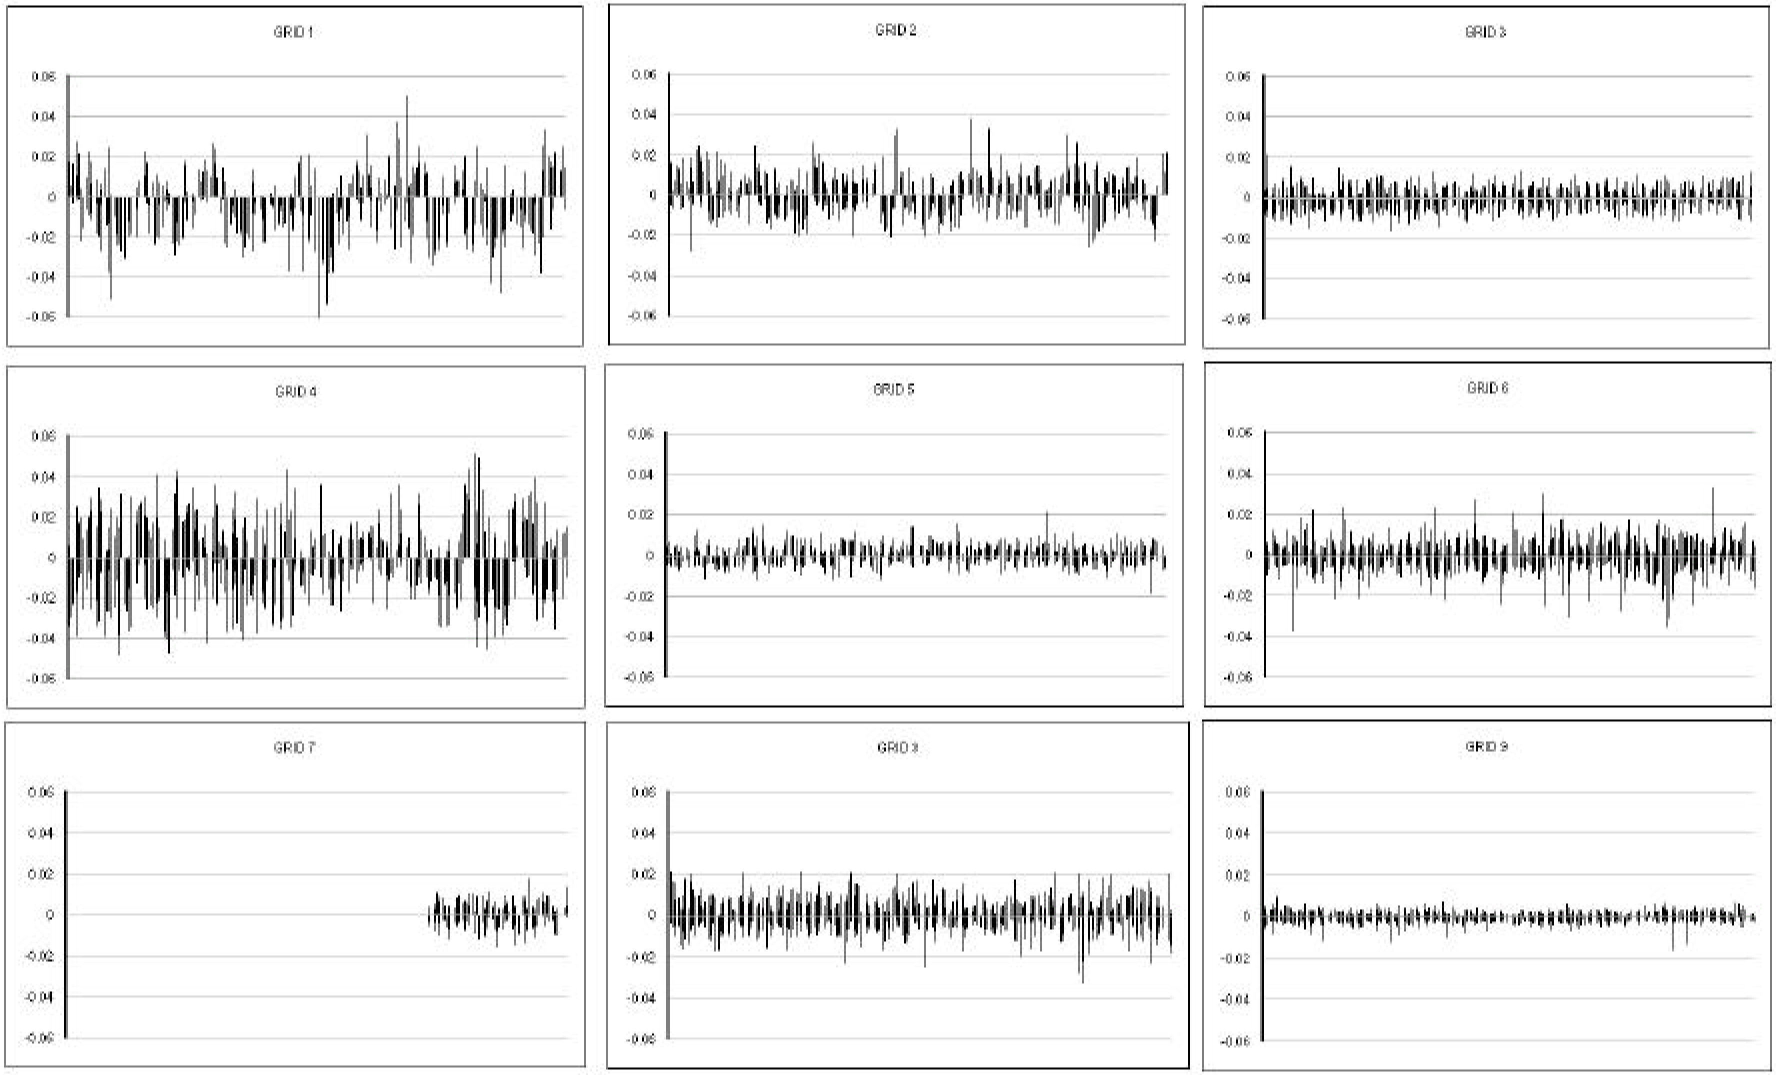

Figure 7

FREQUENCY OF ARBITRAGE

$$\text{ARBITRAGE PROFIT} = \text{MAX}\left[\text{LOG}\left\{\text{BID}^{US} / \text{ASK}^{CN} * \text{ASK}^{FX}\right\},0\right] + \text{MAX}\left[\text{LOG}\left\{\text{BID}^{CN} * \text{BID}^{FX} / \text{ASK}^{US}\right\},0\right]$$

GREY COLUMNS ARE GROSS OF COMMISSIONS.

BLACK COLUMNS ARE NET OF A 2% COMMISSION.