Detecting Pairs Trading

(or any Related Trading)

in Two Financial Securities[*]

Keith Godfrey

The University of Western Australia

12 February 2014

Market-neutral strategies such as pairs trading and merger arbitrage have become increasingly important over the past few decades, yet the literature on them is relatively undeveloped. There are simulation studies which imply theoretical profits, but little empirical work has been done on the actual trading because it is so difficult when trades are anonymous. I describe a technique of "time difference analysis" for studying financial data which is complementary to "time series analysis" and enables empirical inference of pairs trading in two securities. Related trading indicators are computed by analyzing the distribution functions of the time differences between trades, and performance is confirmed by identifying pairs traded from amongst large random sets. Applications of this approach may include searching for traded pairs, inference of dominant trading strategies, day-by-day investigations of merger arbitrage, and studies of rollovers in futures and options.

The traditional approach to studying financial data in two or more securities is to analyze time series of their prices. It is common to compare the performance of two securities by computing their returns at evenly-spaced time intervals. This time series approach has become so ingrained in our thinking that we tend to overlook its imperfections. Forcing any fixed time interval onto irregularly-spaced financial data causes a loss of information or at best a misrepresentation. If

the time spacing is too long there will be many data points discarded in each interval, while if too short there may be few or no data points in each interval, requiring extrapolation. Intraday seasonal (or more correctly diurnal) effects mean different time intervals may suit different times of day, for example Gouriéroux, Jasiak and Le Fol (1999) document more frequent trading activity at the beginning and close compared with slower trading activity in the middle of the day. While early financial studies could contribute with periods of a year, a month, or ultimately a day, today's high frequency trading environment requires an even shorter interval and there is no perfect intraday spacing that fits all financial data.

In this paper I describe a complementary approach which can infer trading in two financial securities by analyzing the time information contained in the trades. It works by investigating the distribution of time differences or relative arrival times of all trade events without considering their prices or volumes. For convenience I refer to this technique in general as "time difference analysis" because its focus on timing rather than prices means it works the opposite way to "time series analysis" which focuses on prices and discards their arrival times. The particular algorithms that I study require no parameterization besides a broad encompassing window, and I will show they can infer the presence of pairs trading and merger arbitrage.

Figure 1 compares and contrasts the concepts of time series analysis and time difference analysis. In time series analysis, each financial time series is constructed from trade or quote prices by an extrapolation rule that discards all but the latest data point in each time interval, ignores the time of its occurrence, and carries this price through any intervals having no data points. Time difference analysis may initially partition the data according to price or volume, then constructs density functions of the differences between the arrival times of the events in the two streams. While time series analysis can help identify security pairs exhibiting price behavior that may be suitable for pairs trading, time difference analysis can help identify security pairs actually being traded.
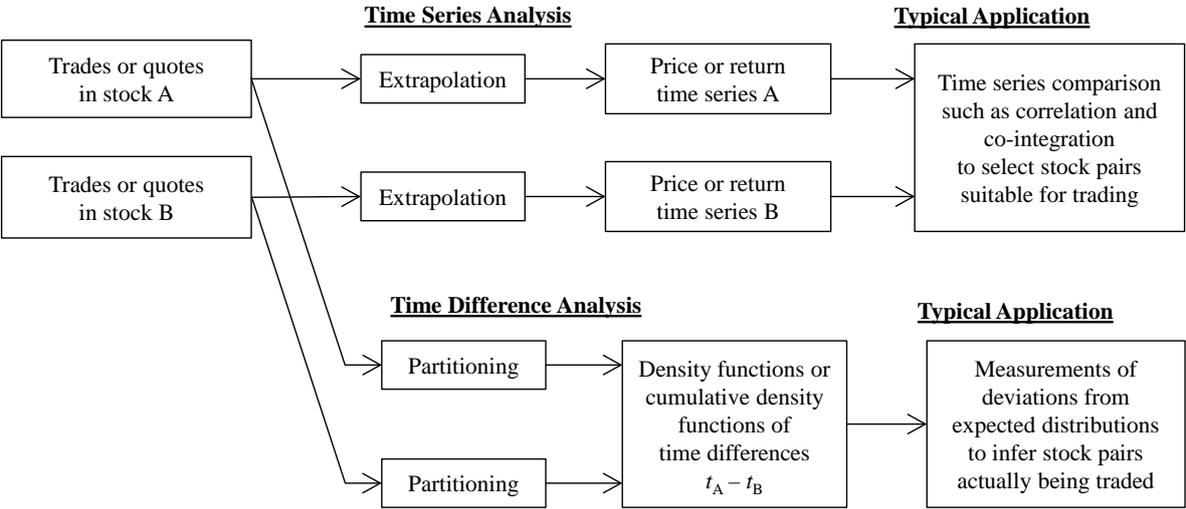
Figure 1: The concepts of "time series analysis" and "time difference analysis". In time series analysis the financial price data is extrapolated to evenly-spaced time intervals discarding all but the last data point in each interval and carrying this forward through empty intervals. In time difference analysis the data may be partitioned based on price or volume, then the differences between trades arrival times in the two securities are aggregated into density functions. While time series analysis can find security pairs with characteristics suitable for pairs trading, time difference analysis can find pairs actually being traded.

This work is by no means the first to look at time differences. In studies of a single security, Engle and Russell (1998) develop a way to analyze irregularly spaced financial events by modeling the duration between the events as a stochastic process: an autoregressive conditional duration (ACD) or a point process with dependent arrival rates. Their model and its extensions surveyed in Pacurar (2006) can be considered as another particular case of time difference analysis. The ACD model is based on event durations which are the time differences between consecutive events. More generally, time difference analysis can compute differences between any two events of which consecutive events are a particular subset. Even more generally, time difference analysis can extend to two securities and study the time differences between events in one security and those in the other, so a single-security study becomes a particular case with both securities the same.

Analysis of two or more streams of financial data concurrently has become increasingly important in the study of multiple-security trading including index tracking, pairs trading, merger arbitrage, and market-neutral strategies. These strategies typically involve opposing long and short positions which require simultaneous trades at the entries and at the exits, and the ideal way to study such trading would be to identify the trades executed by each trader. Unfortunately for researchers, the majority of financial trading data is anonymous. Securities exchanges readily report the dates, times, prices, and volumes traded, but without identifying the traders. As a result it is difficult for researchers to match simultaneous trades in two securities, and even harder to match the subsequent exits with the earlier entries. This article describes a procedure for helping solve the first problem.

The innovation in detecting pairs trading statistically is to compare the density functions of the empirical time differences with that of a uniform distribution. The trades and quotes in any two randomly-chosen unrelated securities will most likely arise from very different trading processes, aside from possible common movements such as basket and program trading. A lack of pairs trading in two securities is then equivalent to their event stochastic processes being practically independent, and I exploit this concept in forming the null hypothesis when detecting related trading. The time differences between all events in one security and all the events in the other should tend toward a uniform distribution when there is no related trading, and deviate from uniform in the presence of related trading.

Rather than employing formal tests of uniformity which give binary decisions without explaining their reasons, I shall utilize graphical methods and relative measures for comparison. This rationale is similar to Diebold, Gunther, and Tay (1998) who use graphical approaches in their evaluations of density forecasts. The additional challenge in detecting pairs trading is there may be a noisy background from other kinds of related trading such as index program trades, which could cause formal tests to reject the null hypothesis too frequently. Graphical techniques, relative measures, and clustering techniques can continue to reveal useful information.

The remainder of this paper is organized as follows. Section 1 reviews studies of trading in multiple securities and the challenges of empirical inference. Section 2 develops the time difference methodology for inferring related trading, initially for detecting any kind of related trading, then for distinguishing particular kinds of related trading such as pairs trading. Section 3 presents results for identifying fundamentally related pairs from pseudo-random sets and finding pairs trading in massive sets formed from index constituent pairs. Section 4 explores the opportunities for further research that are enabled by pairs trading detection.

1. Trading in Two Securities

The idea of analyzing two stocks together for profit has evolved over many years. Long before the era of electronic computers, Livermore (1940) described a method of analyzing two related stocks to determine their common price trend. It was a laborious process and calculated by hand, with the focus on the common movement or "trending" rather than the difference. Half a century later the emergence of computing power helped to automate this pair trending analysis, and it was in this environment the idea of pairs trading evolved. Bookstaber (2007) describes how Gerald Bamberger, a young programmer at Morgan Stanley, started to think of the trending pairs not as a block to be executed but as two sides of a trading strategy. By going long in one and short the other, the net position was market-neutral. Morgan Stanley allowed Bamberger to test his strategy and it made six million dollars in the first year. Nunzio Tartaglia took control of the trading group and reportedly made 50 million dollars for the firm in 1987 (Gatev et al. 2006).

1.1 Mechanism of Pairs Trading

Pairs trading involves the purchase of an underpriced security and the simultaneous sale or short of an over-priced security in such proportions as to maintain a market-neutral position. The combined position is held until either the price difference converges to a target level, or diverges to reach a stop loss. The aim in taking a market-neutral position is to make a profit irrespective of the direction of market movement.

A market-neutral equity strategy in general can involve any number of long and short positions in any combination of securities, provided the overall portfolio has no expected net exposure to risk. Jacobs and Levy (2005) describes several such strategies including market-neutral equity, convertible bond arbitrage, government bond arbitrage, and merger arbitrage, as well as pairs trading. Many hedge funds have thrived from these strategies.

It is also possible to pairs trade two securities from other asset classes such as options and futures, or a mixture such as a stock and an option, or a bond and a future. There is a breadth of literature simulating various kinds of trading strategies across all of these asset classes, but very little empirical evidence of the actual trading.

1.2 Simulations of Pairs Trading Strategies

Simulations can estimate the profitability of trading strategies, and there are many documented studies involving security pairs. Using daily prices from 1990 to 2001, Alexander and Dimitriu (2002) simulate pairs trading among the 30 stocks in the Dow Jones Industrial Average and estimate annual profits of around 10% with 2% volatility and negligible correlation with the market. Hong and Susmel (2003) simulate pairs trading between 64 Asian ADRs and their underlying stocks from 1991 to 2000 and calculate annualized profits of over 33% if investors were to hold the positions for a year. Gatev et al. (2006) simulate pairs trading in U.S. stocks from 1962 to 2002 and calculate annualized excess returns of around 11% before trading costs, or between 2.6% and 4.5% after costs. Chen et al. (2010) conduct a long-run simulation of pairs trading in U.S. stocks using daily and monthly data from 1931 to 2007 and find average returns of 11% to 36% annually before trading costs. The overall message is these strategies could have been successful. Chen et al. (2010) document the returns are diminishing over time, which suggests the market is saturated with pairs trading at the daily level or becoming more efficient.

One way to improve the profitability of a trading algorithm is to trade with a higher frequency using intraday rather than daily data. Large price discrepancies are likely to be short-lived, and Suarez (2005) points out these will be mostly invisible to observers with daily sampling. There is plenty of intraday price data available, and many recent studies use this. Nath (2003) simulates pairs trading in the secondary market for U.S. government debt using trade and quote data from 1994 to 2000 and finds positive excess returns relative to a duration-matched benchmark. Dunis et al. (2010) simulate pairs trading amongst the Eurostoxx 50 index constituents using five-minute prices and conclude pairs trading underperforms the index after trading costs. Bowen et al. (2010) simulate high-frequency pairs trading on a sample of FTSE 100 constituents during 2007 and find the excess returns of the strategy are sensitive to transaction costs, the entry trigger, and delays in execution: a 15-minute delay in execution can eliminate the returns. They also suggest the time of day can be important, noting the majority of returns occur from positions opened in the first hour of trading.

The literature on trading simulations is extensive and growing. Its weakness is it documents only ex-post simulations or paper-trading. The prices used in the simulations are not tradable in practice. Real trading uses ex-ante bid-ask prices and the trades have real market impact. The profitability of real trading is unlikely to agree with the simulations, and astute authors are forthright in acknowledging this limitation.

A more insidious problem in the literature on trading simulations is the conflict of interest between publishing and profiting from such information. Altucher (2004) puts the dilemma bluntly: "If these systems are so good, why not just use them to print money all day long? Why write about them?" Altucher also proposes several valid reasons for publishing, primarily that trading systems evolve continuously, requiring constant research and development. Nonetheless the literature is likely to suffer from a selection bias towards the less profitable or unprofitable algorithms, or a delay in publication while profits are exploited. Morgan Stanley kept silent about its pairs trading strategies evolving in the 1980s, but by the 2000's an abundance of articles had emerged with a common theme that the simulated profits were decreasing.

It would be interesting to know whether arbitrageurs are executing trades similar to those being simulated, how much of such trading goes on, and how closely the profits from real trading match the simulations. These kinds of research question are difficult to answer in an environment where trading is anonymous, and the refereed literature is sparse on such topics.

1.3 Observations of Pairs Trading

There is plenty of anecdotal evidence of pairs trading occurring in practice, but little formal empirical documentation in the refereed literature. Several broad literature searches failed to find any journal articles focused on the detection of pairs trading or the amount and types of trading being undertaken by practitioners. These included searches of databases such as Business Source Premier, JSTOR, ProQuest 5000 International, ScienceDirect, and the Social Sciences Research Network (SSRN). There are papers that mention the topic in passing and others that attempt small-scale tests as part of another study, but studying the actual trading in detail in two or more securities appears to be a space wide open for research.

Anecdotal evidence comes from traders' proud claims of successes, books describing how pairs trading had been conducted, and evidence of government and market responses. Reverre (2001) suggests the arbitrage of Royal Dutch – Shell is a popular model on Wall Street because it has

characteristics close to those of absolute convergence. The trading model assumes the observed value of the price ratio is the superposition of a fundamental function and market noise so they can choose a moving average as an estimator of the fundamental function. Wojcik (2005) describes cases of pair trades which went wrong, implying traders were caught up in those trades at the time. Paul (2008) explains how Australia's share market regulator ASIC relaxed short-selling bans on dual-listed stocks as a result of lobbying from traders, suggesting pairs trades or arbitrage trades were being conducted at the time. In all these reports there is a lack of detail about the types and quantities of trades undertaken. Hedge Fund Research (2011) says merger arbitrage hedge funds have returned on average 1.12% per annum more than the S&P 500 index from 1998 to 2010, but again there are no details of the algorithms employed.

1.4 Empirical Inference of Pairs Trading

Inferring trading between two securities requires a proxy measure for the related trading activity. Do and Faff (2009) mention in one paragraph the possibility of detecting arbitrage activities by examining the spread on the day that follows the opening trigger, arguing that the spread should narrow if a large number of traders follow the prescribed strategy and act on mispricing. This may be one possible way to detect trades but those words were removed in the subsequent journal revision of the paper, perhaps because a narrowing spread can be caused by many other reasons too.

Schultz and Shive (2010) describe in one paragraph how they investigate the trades in dual-class shares from the perspective of studying how prices converge and diverge. They use a process of matching trades as a proxy for arbitrage trades, where a matched trade is defined as a purchase of shares in one class occurring within one minute of a sale of the same number of shares in the other class. They employ an algorithm from Lee and Ready (1991) to classify trades into purchases and sales and investigate the matched sales of expensive shares with purchases of cheap shares, concluding the volume increases when a price differential exists. The result appears to confirm the intuitive proposition that arbitrageurs would exploit such differences. In parallel they find difficulty in explaining why they observe the change in volume from matched trades to be less than the change from single-sided trades, which they proxy by the non-matched trades. They speculate this is due to the single-sided trades being more important in enforcing the prices than round-trip arbitrage trades, a concept described earlier as "one-way arbitrage" by Deardorff (1979). An alternate explanation could be that any experiment testing whether the volume of single-sided trading exceeds that of round-trip trading involves a joint hypothesis test with the choices of proxies. The choice of matching only on equal volumes traded within one minute in opposite inferred directions of aggression is likely to be a sub-optimal proxy for the single-sided arbitrage trades.

The anonymity of trades makes it difficult to relate trades in two securities. Specific challenges include:

- An algorithm can infer the direction of aggression in one security with reasonable success, but it is harder to match the orders across two securities. Trading in two securities can involve either a market or limit order in the first security followed by a market order in the other. The two choices for the initial order will be inferred empirically as having opposite directions of aggression, leading to a risk of misclassification;

- Securities exchanges tend to split orders into smaller lots during execution and matching, so any algorithm based on traded volume will be noisy with false positives and negatives. This leads to a risk of misclassification of paired trades being unrelated single-sided trades, and vice versa; and

- Introducing any fixed-size time window onto an exact-matching process will introduce artifacts which vary with window size. As an example, changing the fixed 1-minute window in Schultz and Shive (2010) from 1 minute to either 30 seconds or 2 minutes is very likely to change the results dramatically.

The abundant streams of financial data deserve more general ways of inferring related trading, approaches that can be forgiving of the inherent uncertainty in classifying the observed trades. As financial researchers we also need to exercise caution in interpreting the results from using such proxies and measures, devising tests to inspire confidence in their accuracy.

2. Empirical Inference of Related Trading by Time Difference Analysis

Studying related trading in two securities is difficult because traders have a financial incentive to keep profitable strategies confidential, and securities exchanges collaborate by concealing traders' identities. Researchers have to live with the uncertainty of whether any two trades are part of the same strategy from the same trader. Attempts to match such trades without proper identification will inevitably be statistical and will likely need large quantities of trading data to provide sufficient statistical confidence. This section develops one such approach.

2.1 Empirical Inference of Related Trading based on Times between Trades

There are two reasons for trading in two or more securities simultaneously: either it is to capture a joint trend as in Livermore (1940), or it is to exploit a difference and maintain market neutrality as in Gatev et al (2009). Trades in the same direction can include program trades where baskets of securities are purchased or sold simultaneously. Trades in opposite directions can include strategies such as pairs trading, index arbitrage, and merger arbitrage. All of these trading rules, whether in the same or opposite directions, represent related trading in the securities. The task of detecting any kind of related trading, regardless of its type, comes ahead of discerning the strategies into different types.

The analysis here focuses on two securities, any two securities, irrespective of fundamental relations or observed price cointegration characteristics. The approach makes use of only the time differences between trades, without influence from the price or volume. Additional information from price, volume, and subsequent inferences of the directions of trades, can be used in partitioning the data beforehand to infer different types of related trading.

Consider two securities AA and BB which have sets of trades

$$\mathbf{AA} = \{ AA_1, AA_2, AA_3, \ldots, AA_{\|\mathbf{AA}\|} \}$$

$$\mathbf{BB} = \{ BB_1, BB_2, BB_3, \ldots, BB_{\|\mathbf{BB}\|} \}$$

(1)

If we denote the signed time difference between trades AA_i and BB_j as $\mathbf{\Delta t}(AA_i, BB_j)$ we can define the set of all such time differences by $\mathbf{\Delta t}(\mathbf{AA} \times \mathbf{BB})$. This is the set of time differences to be analyzed, and is the most general set for time difference analysis, but it is likely to be too

large for practical operations. Its size is that of the Cartesian product set formed by matching every trade in **AA** with every trade in **BB**:

$$\left\| \mathbf{\Delta t}(\mathbf{AA} \times \mathbf{BB}) \right\| \;=\; \left\| \mathbf{AA} \times \mathbf{BB} \right\| \;=\; \left\| \mathbf{AA} \right\| \times \left\| \mathbf{BB} \right\| \tag{2}$$

To enable practical computation we can define a subset in which the time differences between trades is limited to a domain $\left[-T, T \right]$ for suitable choice of T and denoted $\mathbf{AA} \times \mathbf{BB} : |\mathbf{\Delta t}| \leq T$. The precise value of parameter T is not critical. The only requirement is it must be sufficiently large to capture the time differences that are of interest.

We can similarly define a subset for any interval of time differences $\left[T_1, T_2 \right]$ as

$$\mathbf{AA} \times \mathbf{BB} : T_1 \leq \mathbf{\Delta t} \leq T_2 \;=\; \left\{ \begin{array}{l} (AA_i, BB_j) \;:\; AA_i \in \mathbf{AA}, \; BB_i \in \mathbf{BB}, \\ \qquad\qquad T_1 \leq \mathbf{\Delta t}(AA_i, BB_j) \leq T_2 \end{array} \right\} \tag{3}$$

This notation enables us to define an empirical measure or relative frequency of time differences based on the interval $\left[T_1, T_2 \right]$ with $-T \leq T_1 \leq T_2 \leq T$ as

$$
\begin{aligned}
\hat{P}_{\mathbf{AA} \times \mathbf{BB} : |\mathbf{\Delta t}| \leq T} \left[T_1 \leq \Delta t \leq T_2 \right] \;&=\; \frac{\text{number of products with } \Delta t \in \left[T_1, T_2 \right]}{\text{number of products with } \Delta t \in \left[-T, T \right]} \\[2mm]
&=\; \frac{\left\| \mathbf{AA} \times \mathbf{BB} : T_1 \leq \mathbf{\Delta t} \leq T_2 \right\|}{\left\| \mathbf{AA} \times \mathbf{BB} : -T \leq \mathbf{\Delta t} \leq T \right\|}
\end{aligned}
\tag{4}
$$

The empirical distribution or cumulative density function (CDF) is

$$
\begin{aligned}
\hat{F}_{\mathbf{AA} \times \mathbf{BB} : |\mathbf{\Delta t}| \leq T} (t) \;&=\; \hat{P}_{\mathbf{AA} \times \mathbf{BB} : |\mathbf{\Delta t}| \leq T} \left[-T \leq \Delta t \leq t \right] \\[4mm]
&=\; \frac{\text{number of products with } \Delta t \in \left[-T, t \right]}{\text{number of products with } \Delta t \in \left[-T, T \right]} \\[4mm]
&=\; \frac{\left\| \mathbf{AA} \times \mathbf{BB} : -T \leq \Delta t \leq t \right\|}{\left\| \mathbf{AA} \times \mathbf{BB} : -T \leq \Delta t \leq T \right\|}
\end{aligned}
\tag{5}
$$

These measures are readily calculable from the records of anonymously reported trades by counting the number of pairs having time differences within each range, although this may involve a large number of computations.

The insight to detecting related trading between two securities is to recognize the following:

Proposition 1: **If the trading in two securities is unrelated, and the opportunity to trade is available continuously, the time differences between trades in the two securities should approximate a uniform distribution.**

This idea is illustrated in panel (a) of Figure 2. When there is no related trading between the two securities, the time differences between any two trades selected at random from the two continuously-traded securities should approximate a uniform distribution on $\left[-T, T \right]$.

Alternatively, the presence of related trading should distort the distribution according to the time differences of the paired trades executed.

(a) Empirical CDF of time differences between trades

Uniform distribution

possibly not 0.5

(b) Deviation of CDF from uniform

possibly not 0

(c) Empirical CDFs of trade time differences each way

Uniform   Uniform

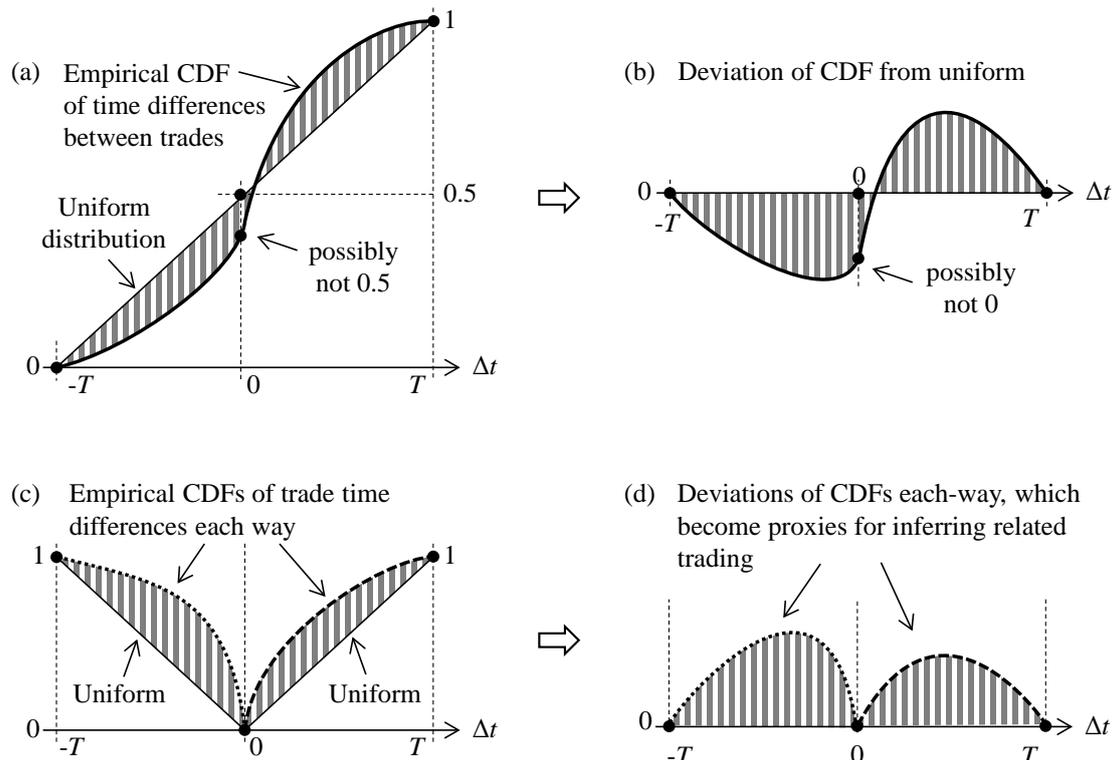(d) Deviations of CDFs each-way, which become proxies for inferring related trading

Figure 2 Concept of inferring related trading using the cumulative density functions (CDFs) of time differences between trades in two securities. The horizontal axis in each panel represents the time difference from any trade reported in one security to all trades reported in the other security within the limits $[-T, T]$. Panel (a) illustrates an empirical distribution of these time differences compared with a uniform distribution which is anticipated when no related trading occurs. The grey shaded area (the difference between the empirical CDF and the uniform distribution) is a proxy for the amount of related trading between the securities. Panel (b) shows this more explicitly by computing the difference. The limitation is the zero-crossing may not be at $\Delta t=0$ if there is an excess of trades where one stock leads the other. Panels (c) and (d) illustrate computations of the CDFs separately for $[0, -T]$ and $[0, T]$. The grey regions in panel (d) are proxies for the amount of related trading inferred in each direction. These regions can be measured by the Cramér von Mises criterion and the Kolmogorov-Smirnov distance to give numerical indicators of inferred related trading.

It is important to remember this procedure is analyzing all time differences from each trade in one security to all the trades in the other security, not just to the nearest or consecutive trade. Studying the distribution of the *nearest* arrival time rather than to *all* arrival times would lead to Poisson distributions with the problem of determining parameters which would be different for each security. Studying time differences between all pairs of arrival times offers a simpler comparison with a uniform distribution which does not require a parameter.

The assumption of trading being available continuously contrasts with the reality that most stock exchanges trade for fixed periods each day. For example The NYSE opens from 9:30 to 16:00 daily, giving a window of 6.5 hours for trading. Fortunately the null hypothesis in that situation still looks very similar to a uniform distribution when studying the small time differences between trades in pairs trading. Consider an exchange that reports trades with millisecond resolution and is open for a window of W milliseconds per day, which would be $W = 23,400,000$ for the NYSE. If two trades occur in the same day, there are W ways they can occur on the same millisecond, $W-1$ ways they can be 1 millisecond apart, $W-2$ ways they

can be 2 milliseconds apart, and so on down to the final 1 way they can be $W-1$ milliseconds apart. This linearly decreasing density function means the cumulative density function on $[0,T]$ is a decreasing square law of the form

$$F_{[0,T]}^{**}(t) \quad = \quad \frac{Wt - t^2/2}{WT - T^2/2} \tag{6}$$

rather than a uniform distribution. Importantly this is flat around $t = 0$ and especially for large T. If we consider a limit of say $T = 10,000$ microseconds which is just 0.043% of the NYSE daily opening window of 6.5 hours, this approximates a cumulative uniform distribution

$$F_{[0,T]}^{*}(t) \quad = \quad \frac{t}{T} \tag{7}$$

and this means we can continue to work with Proposition 1.

The uniform empirical measure over the interval $t \in [-T,T]$ has CDF

$$F^{*}(t) \quad = \quad \frac{t+T}{2T} \tag{8}$$

so the observed CDF of time differences will differ from the uniform distribution by

$$[\hat{F} \text{-} F^{*}](t) \quad = \quad \hat{F}_{\mathbf{AA \times BB}:\, |\Delta t| \leq T}(t) \quad - \quad F^{*}(t) \tag{9}$$

which is the curve illustrated in panel (b) of Figure 2.

To infer a measure of related trading, we want a single metric to capture the difference: a distance measure between the empirical CDF and the uniform distribution.

Proposition 2: **A distance measure between the empirical CDF of time differences between trades in two securities and a uniform distribution can be a proxy for the amount of related trading.**

Common alternatives for this distance measure include the Cramér von Mises (CVM) criterion

$$\omega^2 = \frac{1}{2T} \int_{-T}^{T} \left(\hat{F}(t) - F*(t) \right)^2 dt \tag{10}$$

and the Kolmogorov-Smirnov (KS) distance

$$D = \sup_{t \in [-T,T]} \left| \hat{F}(t) - F*(t) \right| \tag{11}$$

I call the single numerical result from any such measure a Related Trading Indicator (RTI) and I use subscripts such as CVM or KS to denote the chosen measure. The RTI computations for the CVM and KS measures are summarized as

$$\text{RTI}_{\text{CVM}} \quad = \quad \sqrt{\frac{1}{2T} \int_{-T}^{T} \left(\hat{F}_{\mathbf{AA \times BB}:\, |\Delta t| \leq T}(t) \quad - \quad F^{*}(t) \right)^2}$$

$$\tag{12}$$

$$\text{RTI}_{\text{KS}} \quad = \quad \sup_{t \in [-T,T]} \left| \hat{F}_{\mathbf{AA \times BB}:\, |\Delta t| \leq T}(t) \quad - \quad F^{*}(t) \right|$$

Equation (12) specifies the continuous-time form of the RTIs. In practice the indicators are to be computed from time differences between reported trades, which by their nature are resolved to the discrete time intervals as the trade reporting. Denoting the time resolution by τ (which is typically at most one microsecond in modern financial data recording) and rewriting the maximum time difference of interest as $T = N\tau$ for some integer N, the computation becomes

$$\text{RTI}_{\text{CVM}} = \sqrt{\frac{1}{2N+1} \sum_{n=-N}^{N} \left(\hat{F}_{\mathbf{AA} \times \mathbf{BB}: |\mathbf{\Delta t}| \leq T}(n) - F^{*}(n) \right)^2}$$

(13)

$$\text{RTI}_{\text{KS}} = \sup_{n \in [-N, N]} \left| \hat{F}_{\mathbf{AA} \times \mathbf{BB}: |\mathbf{\Delta t}| \leq T}(n) - F^{*}(n) \right|$$

This is the discrete-time form of the related trading indicators. The RTIs are readily calculable from Equations (5) and (13) although the number of computations involved may be large.

It should be possible to use either the Cramér von Mises or Kolmogorov-Smirnov forms of the RTIs to measure inferred related trading, although there may be subtle differences between them depending on the characteristics of the security pairs being analyzed. The CVM criterion computes a root-mean-square of the grey shaded regions in Figure 2, while the KS distance measures their maximum deviations. We can imagine the CVM criterion may be better at inferring related trading when a security pair is traded by many traders having a wide range of time differences in their execution strategies, while the KS distance may be better in pairs where trading is dominated by a single trader executing in a narrow range of time differences.

Analyzing the empirical distributions of the time differences means we do not require prior knowledge of the many trading strategies that may be present in the data. We can expect any kind of related trading to show an excess relative frequency of trades occurring at particular time differences, and traders will want to keep those time differences as short as possible. With modern financial data, a time difference limit of $T = 10$ seconds is likely to be sufficient to catch the majority of trades of interest.

The use of numerical distance measures rather than binary tests of uniformness means the relative trading detection is relative rather than absolute. Such approaches can distinguish pairs traded from amongst large sets of stock pairs with background related trading noise arising from the fact many of them may be traded in baskets for index purchases and sales. Security pairs with larger indications can be interpreted to have more related trading than pairs with lower indications, ceteris paribus, although the actual value or magnitude has no physical meaning. The related trading indicators make no assumptions about the underlying shape of the common distributions.

The related trading indicators can accommodate the varying arrival rates of trades from intraday seasonal (or more correctly diurnal) effects because they are looking at the difference in times in trades between the two securities. The amount of time taken between executions of a trader's pair of orders is likely to remain the same regardless of whether the trading environment is busy or quiet. On the other hand, these intraday variations in trading frequency will cause a greater spreading of time differences that do not arising from related trading. The effect overall is likely to help rather than hinder the acuity of time difference analysis.

2.2 Distinguishing Aggressive Pairs Trading, Passive Pairs Trading, and Program Trading

The methodology described thus far is designed to infer any related trading between two securities and is based on the time differences alone. The information from prices and volumes traded has not yet been utilized. The way to apply the RTI methodology to more complex tasks is to recognize it can be applied to any subsets of trades, and partition the trades appropriately.

Proposition 3:   **Related trading can be inferred between subsets of trades in two securities that are inferred or constructed from price, volume, and time information to distinguish particular types of related trading.**

Aggressive pairs trading can be inferred from paired sets of trades having opposite directions of aggression. Passive pairs trading and program trading can be inferred from paired sets of trades having the same directions of aggression. Program trading can be distinguished from passive pairs trading by a shorter time difference between market orders.

The idea is to partition the set of trades $\mathbf{AA} \times \mathbf{BB} : |\Delta t| \leq T$ into subsets of interest such as those having particular inferred combinations of purchases and sales, or those occurring at times of inferred entries and exits of particular trading algorithms. The choice of partitions determines the information to be deduced. We can begin incorporating price information to distinguish aggressive and passive pairs trading. The aggressor in each trade can be inferred from the price of the trade relative to the prevailing bid-ask spread, most simply by bisection, or by more complicated algorithms such as those in Lee and Ready (1991). The simple bisection approach is sufficient here to show the capabilities of the methodology. Alternatives and enhancements can be trialed later. From here on, trades occurring above the midpoint of the prevailing bid-ask spread are classified as initiated by the buyer, and those below the midpoint are classified as initiated by the seller.

Analyzing the relative aggressions of any two trades becomes complicated. Opening or closing a pairs position will involve trades in the two securities but the style can be either aggressive or passive. An aggressive pairs trade means submitting a pair of market orders in opposite directions, while a passive pairs trade involves waiting for a limit order to be filled in one security before executing a market order in the other. The first situation may be inferred from the trading logs as a pair of trades in opposite directions. The second may be inferred from the logs as two trades in the same direction because the direction of aggression inferred from the trade with the limit order is the opposite direction to the passive limit order. Program or basket trades may also be inferred with aggressions in the same direction because they arise from market orders in the same direction, and this leads to potential confusion in distinguishing passive pairs trading from program or basket trading. One possible solution is to partition by time difference as well as by relative direction of aggression. The trades executed from the simultaneous market orders of a program trade should occur closer in time than those requiring a trader to react to news of a passive order being filled before ordering the second trade.

Figure 3 illustrates the set of trade combinations $\mathbf{AA} \times \mathbf{BB}$ being reduced initially to the subset $\mathbf{AA} \times \mathbf{BB} : |\Delta t| \leq T$ then subdivided further by time and relative directions of trade. When a related trading indicator (RTI) is calculated for one of these subsets, it can be called an aggressive pairs trading indicator (APTI), passive pairs trading indicator (PPTI), or basket trading indicator (BTI). In each case the calculation is the same for the general RTI but applied to the particular

subset of trades. The APTI, PPTI, or BTI can be suffixed similarly with a subscript for the distance measure employed such as CVM or KS.
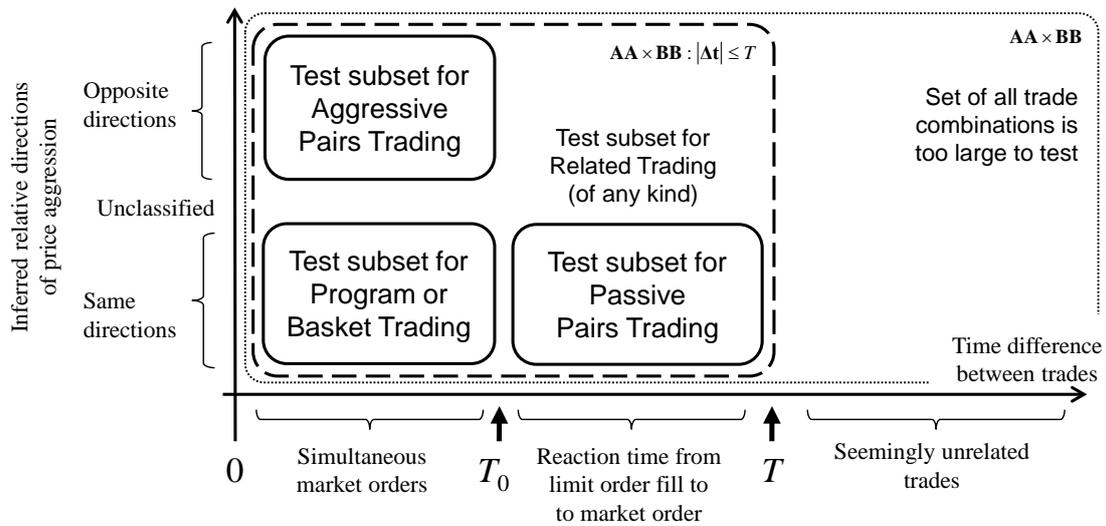

Figure 3   Classification of trading from two securities into subsets based on the inferred relative directions of aggression and the absolute time difference between the trades. Trades in each security are classified as buy or sell according to the trade price and the prevailing bid-ask spread. Trade pairs are then classified in two dimensions with the vertical axis separating those having opposite directions of aggression (buy sell or sell buy) from those with the same directions (buy buy or sell sell), and the horizontal axis being the absolute time difference between the trades.

3. Empirical Explorations of Related Trading

3.1 Validation of the Methodology

The accuracy of the related trading indicator approach can be verified if it can identify security pairs in which pairs trading is anticipated, relative to randomly selected pairs in which there is no such prior expectation. The NYSE contains several pairs of closely-related securities which are either American stocks or American Depository Receipts (ADRs). Table 1 shows the four pairs selected here to verify the related trading methodology.

Tickers	Securities	Description
RDSa and RDSb	Twin ADRs of Royal Dutch Shell Plc	Shell is a global petrochemical company with two share classes listed on the London Stock Exchange. Each has an ADR on the NYSE.
UN and UL	Twin ADRs of Unilever NV and Plc	Unilever is a global manufacturer of nutrition and hygiene products. It has a dual-listed Dutch-British company structure, both having ADRs.
CCL and CUK	Twin stock and ADR of Carnival Corporation	Carnival is a global cruise line and holiday provider registered in Panama and listed on the NYSE that took over a UK company P&O Cruises in 2003 which has an ADR.
BHP and BBL	Twin ADRs of BHP Billiton Ltd and Plc	BHP Billiton is the world's largest diversified resources group, with a dual-listed Australian-British structure formed in 2001. Both have ADRs.

Table 1   The four pairs of closely-related securities chosen for validation of the methodology.

The four pairs selected are Royal Dutch Shell, Unilever, Carnival, and BHP Billiton. These are all economically-significant entities and each has twin securities that can be expected to be

traded together in strategies such as pairs trading. Royal Dutch Shell was studied extensively by Rosenthal and Young (1990) and Froot and Dabora (1999) when Shell had a dual-company structure. In 2005 Shell restructured into a single British parent company having two share classes derived from Dutch and British origins. The shares nowadays have equal intrinsic value instead of the 60:40 split in favor of the Dutch prior to the restructure.

To test the methodology we need to compare these candidates with similar pairs not expected to be involved in pairs trading. Comparison sets from 2008 through 2010 are constructed from pairs of securities having similar trading characteristics to each of the main pairs. Picking the four closest securities to RDSa by number of trades in the year, plus four more by trading volume, and four by dollar volume, gives a set of 13 securities with characteristics similar to RDSa (including RDSa itself). Building a set similarly of 13 securities around RDSb enables a set of $13 \times 13 = 169$ security pairs to be constructed from all combinations of the RDSa and RDSb sets, of which one pair is expected to show pairs trading and 168 are the pseudo-random comparison pairs. This approach is repeated for each of the four pairs of interest and for each time period to be tested.

Trading data are obtained from Thomson Reuters Tick History with thanks to SIRCA for access. In 2008 there were 54 009 328 trades in the 13 securities in the set constructed around RDSa and 16 063 792 in the set around RDSb. Pairing all these trades gives a Cartesian product set of size 822 418 826 468 768 which is unwieldy for analysis. Limiting the time difference between trades to $T = 100$ seconds reduces the number of trade pairs for analysis to 37 305 542 257. This is still a large set but it is manageable.
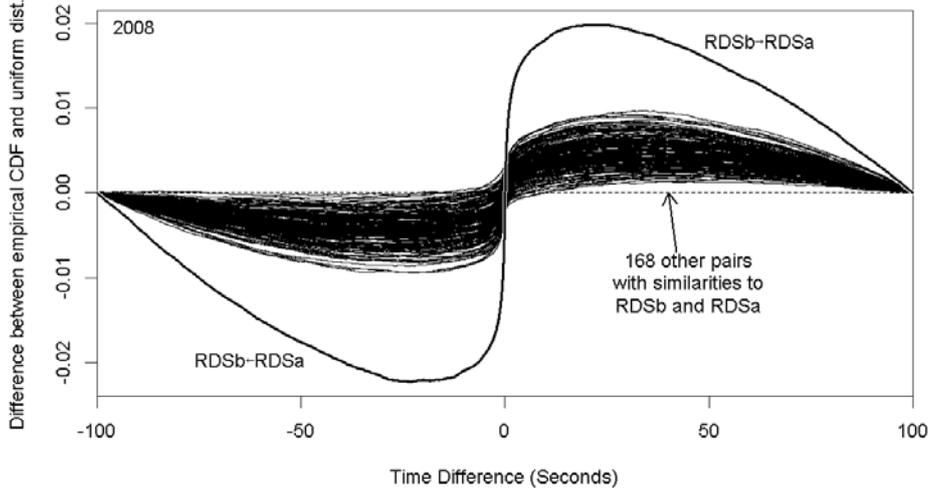

Figure 4 Deviations from the uniform distribution of the CDFs of time differences between trades in Royal Dutch Shell pair RDSb-RDSa and 168 pseudo-random pairs in 2008. The set of pairs is constructed by combinations of RDSb and 12 similarly-traded securities with RDSa and 12 similarly-traded securities. The curve for RDSb-RDSa stands out from all the other pairs by a significant distance, indicating a much greater presence of related trading. Repeating this test for Unilever, Carnival and BHP Billiton in three disjoint time periods validates the methodology of computing numerical related trading indicators (RTIs) by measurement of these curves.

Figure 4 shows the empirical distributions or CDFs of the 169 pairs in the Royal Dutch Shell test set in 2008 computed from Equation (5). One pair stands out a long way from the other 168 and it is RDSb-RDSa as predicted. A similar result occurs in the pairs for Unilever, Carnival, and BHP Billiton. The results for all four pairs are robust to calendar year when repeated for sets generated separately for 2009 and 2010. It is unlikely these results are obtained by chance for all

14

four pairs and all four years. The 37 305 542 257 trade combinations analyzed for the Royal Dutch Shell set in 2008 are independent of the 43 696 415 909 combinations in 2009 and the 32 688 816 870 in 2010, and again from the Unilever, Carnival and BHP Billiton sets in 2008 to 2010 which range from 17 954 524 703 to 171 083 917 733 combinations. These numbers are large and come from 12 disjoint sets. The methodology appears to be working for these securities.

Besides validating the methodology, this set of tests also discovered a few pairs in the pseudo-random comparison sets that appear to be traded. Figure 5 shows the RTIs calculated for the 169 security pairs around Unilever (UN and UL) in 2009. The UL-UN pair stands out by a long distance from the main clump of pseudo-random pairs, verifying the methodology, but there is an additional feature of interest. The pair SWK-BDK also stands out albeit by a smaller margin. This pair comprises Stanley Works and Black & Decker, two competing manufacturers of tools and hardware which merged on 12 March 2010 to form the Stanley Black and Decker Corporation. The empirical inference of related trading here appears to have found merger arbitrage in the pre-merger securities (the 2009 data is at least two months prior to the merger date). This deserves further study. The ability to infer such related trading may open more avenues for empirical research into mergers and acquisitions, including empirical inference of possible insider trading activity ahead of particular announcement dates.
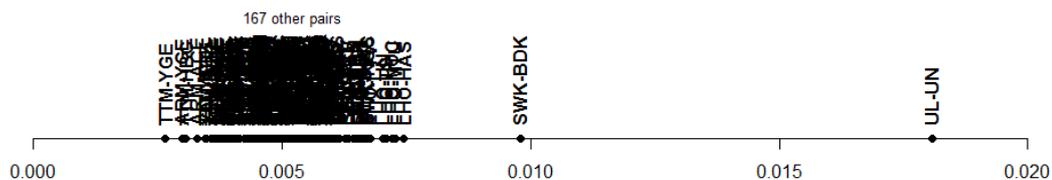

Figure 5 Related Trading Indicator RTI $_{CVM}$ for the Unilever pair UL-UN and 168 pseudo-random pairs in 2009. The set of pairs is constructed by all combinations of UL and 12 similarly-traded securities with UN and 12 similarly-traded securities. The related trading methodology works as anticipated by distinguishing the UL-UN pair from the others, but a second pair SWK-BDK also stands out to a lesser extent. This pair is Stanley Works and Black & Decker, two competing manufacturers of tools and hardware which merged on 12 March 2010 to form the Stanley Black and Decker Corporation. Related trading (in this case merger arbitrage) is inferred in those securities in the 2009 data more than two months ahead of the merger.

A further test on the BHP Billiton data from 2012 finds the BHP-BBL pair stands out significantly compares with pairing of either BHP or BBL with the other securities in the NYSE. Figure 6 shows the CDF differences for BHP mapped against 1,696 other NYSE securities in 2012. The CDF of the BHP-BBL pair differs significantly more from the uniform distribution than all the other pairs, as predicted. It is interesting to observe the left side of the BHP-BBL curve (which corresponds to BBL being traded before BHP) shows greater deviation from uniform than the right side (which corresponds to BHP being traded first). This could arise from traders who use limit-market order pairs choosing to place the limit order in the less frequently traded stock (BBL) to earn the bid-ask spread on that stock before placing a market order in the more liquid stock (BHP).
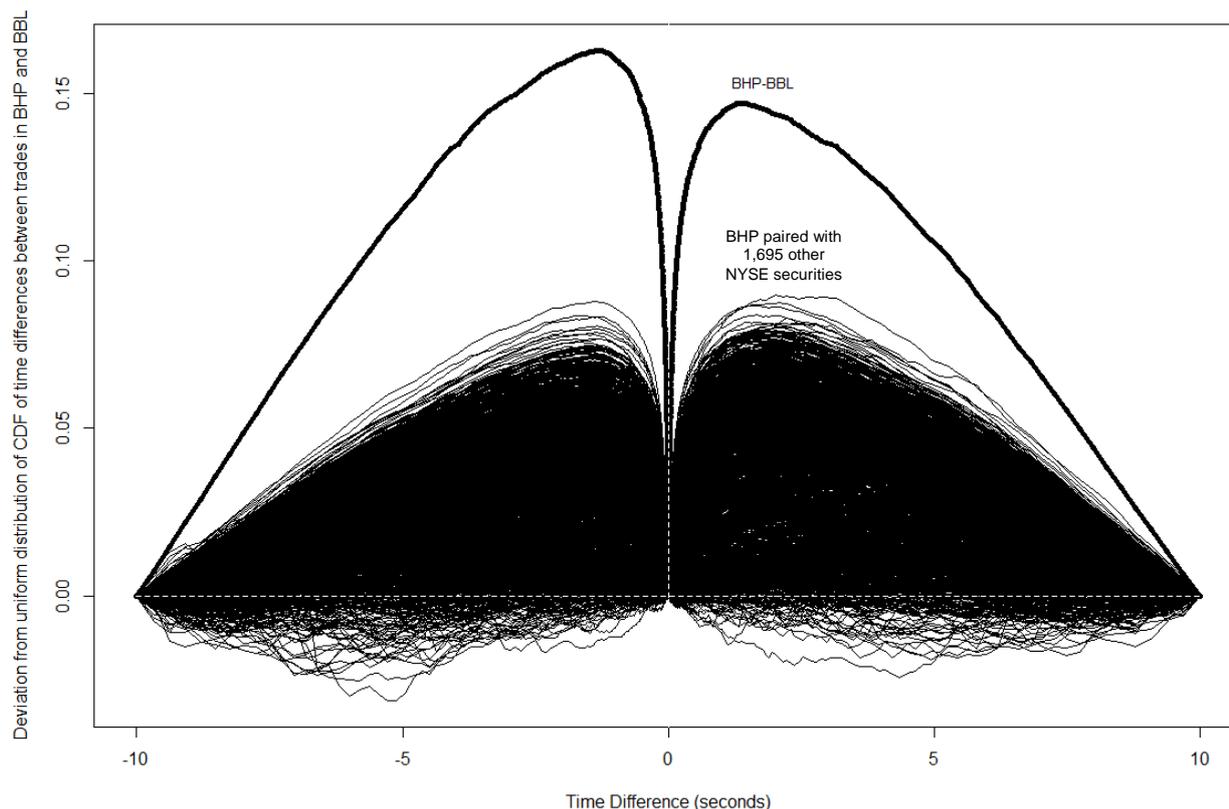
Figure 6 Deviations of CDFs of time differences between trades in BHP and trades in 1,696 other NYSE securities most-frequently traded, one of which is BBL. The BHP-BBL pair stands out, confirming the techniques described are able to distinguish fundamentally related pairs trading from the background noise of other related trading. The large gap below BHP-BBL favors graphical inspection and clustering techniques rather than format tests which would reject a null hypothesis in many cases. This graph shows there is related trading between BHP and several hundred securities, but BBL stands out by a very large margin, as predicted by the dual company relationship. A further observation on the BHP-BBL curve is the left side is taller than the right side, which means BBL is traded more frequently as the first stock in each pairs trade. This may be because traders who use limit orders in the first trade would rather wait for the less-liquid stock trade first, in this case BBL.

Testing robustness to intraday seasonal (or rather diurnal) variations in trading activity involves partitioning the data within each day before computing the CDFs and RTIs on each of the partitioned data sets. Figure 7 shows the results for six hourly partitions each day from 10am to the 4pm closing time of the NYSE, throughout the year 2012. The data points are plotted in two dimensions with the frequency of trading along the horizontal axis to further distinguish the security pairs, although in this example it can be seen the RTIs on the vertical axis are capable by themselves of identifying the BHP-BBL pair. The six data points for BHP-BBL in Figure 7 plot very closely together despite diurnal variations in the overall trading frequency. This suggests that the amount of trading within 10 seconds (the horizontal dimension) and the amount of inferred related trading (the vertical dimension) each remain roughly constant throughout each day, irrespective of the overall trading frequency.

Additional testing has verified the related trading indicator methods are robust to the choices of securities exchanges on which the trades have executed. This is important because differences in the trade reporting mechanisms between exchanges could distort the time difference analysis. The NYSE listed securities can typically be traded on several exchanges and electronic communication networks besides the NYSE itself, and it can be expected that each such

platform will each require slightly different amounts of time to report their trades. To test robustness to this, the BHP-BBL trade data sets from each of the years 2008 to 2010 were partitioned by each of the 25 combinations of the five most frequently traded exchanges. In all cases the RTI approach continued to identify the BBL-BHP pair correctly regardless of the exchange pair and whichever security traded first. Asymmetrical patterns emerged in the graphical plots which imply timing differences do exist between exchanges, and these were insufficient to alter the results.
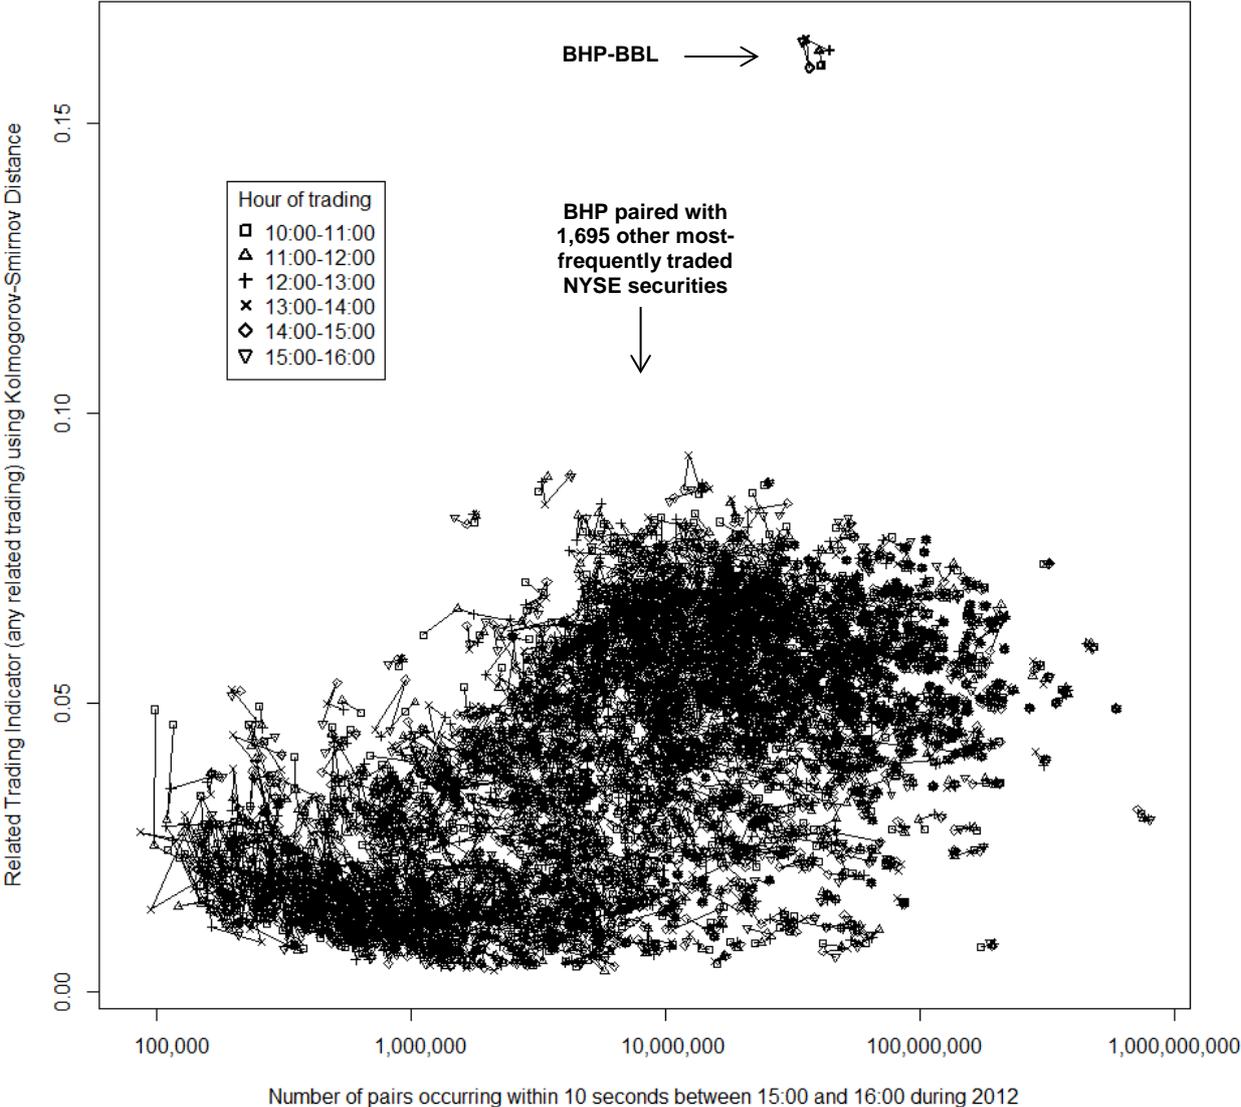

Figure 7 Test of robustness to intraday seasonal (or diurnal) effects. The trading data from 2012 for BHP and 1,696 other most-frequently traded NYSE-listed securities (one of which is BBL) is partitioned by hour of day into six sets corresponding to 10am-11am, 11am-12pm, through to 3pm-4pm (15:00-16:00). Time difference analysis techniques are applied on each set independently to compute related trading indicators, which are plotted as a scatter diagram. The horizontal axis is the number of trades in the 3pm-4pm timeslot as a proxy for the frequency of trading, and the vertical axis is the related trading indicator. The six timeslot subsets for each security pair are joined by lines to show visually how they each tend to stay roughly in the same part of the trading versus RTI space. The six timeslots for BHP-BBL are seen to plot very closely together, confirming the time difference analysis technique is robust to diurnal effects, and all at a long distance from all the other pairs and timeslots, verifying once again that the methodology works.

17

3.2 Exploration for Pairs Trading amongst Index Constituents

With more than 40 000 equity securities listed on exchanges around the world, there is a potential for related trading in more than 800 million security pairs. In practice pairs trading can be performed only when one or both of securities can be shorted so we can restrict the universe to pairs of short-sellable securities. Interactive Brokers lists 15 383 of these in March 2013, so the practical universe still contains more than 118 million pairs. Random searches for pairs trading in this space are unlikely to have success. The explorations here shall be restricted to four sets of index constituents: the S&P 500, NASDAQ 100, S&P MidCap 400, and FTSE 100. The numbers of trade combinations to analyze in these sets are still large, as listed in Table 2.

Figure 8 shows the result of computing the aggressive pairs trading indicator $APTI_{CVM}$ on the pairs of index constituents from the S&P 500 index with time limit T=10 seconds. The S&P 500 data set was obtained from Thomson Reuters Tick History using chain 0#.spx which returned 513 codes including those added, deleted, and changed during the period as well as additional codes for distinct classes of shares. The 1 450 933 036 trades reported in the first six months of 2011 create a Cartesian product set of size 1 048 326 510 036 240 846 from the 131 328 security pairs. With the time difference limited to 100 seconds, the number of trade products to analyze is reduced to 87 529 904 232 756. This is still a large number and took a month to process on an eight-core laptop. The analysis was also performed with limits of 1, 2, 5, 10, 20 and 50 seconds for robustness. We can envisage improvements in performance over time as computer hardware evolves because the methodology developed here and experimental sequence are both entirely parallelizable.

Index	Chain RIC	Time period	Codes	Code pairs	Trade products within T=100 seconds
S&P 500	0#.spx	2011 H1	513	131 328	87 529 904 232 756
NASDAQ 100	0#.ndx	2010	107	5 671	15 127 152 881 204
S&P MidCap 400	0#.mid	2010	451	101 475	3 637 794 060 531
FTSE 100	0#.ftse	2010	117	6 786	151 046 067 927

Table 2 The four equity indices explored for pairs trading among constituents. RIC means Reuters Instrument Code in Thomson Reuters Tick History, and Chain RICs listed here provide the constituent lists of each underlying indices. The periods of analysis are the 12 months of 2010 except for the S&P 500 which is analyzed during the first six months of 2011. The S&P period was reduced to six months because there are many more trades amongst its constituents than in all the other indices combined, and for robustness its period was chosen not to overlap with 2010. The numbers of codes analyzed in each case exceeds the size of the index because of additions and deletions during the periods and the presence of additional codes for different classes of shares. In general for N codes there are $(N^2-N)/2$ pairs. Even after reducing the trade combinations to those executed within T=100 seconds, the numbers of trade combinations for analysis remain large.

The pairs from the S&P 500 index in Figure 8 are plotted in two dimensions with the aggressive pairs trading indicator $APTI_{CVM}$ on the vertical axis and the number of trades within time difference of T=10 seconds on the horizontal axis. The number of pairs plotted is 127 260 after discarding those securities which trade less than 1% as frequently as the most traded security. The diagram can be described as having a few pairs such as CTL-Q and AYE-FE standing out by a wide margin, and a diagonal frontier of pairs with the most inferred related trading facing the upper right corner.
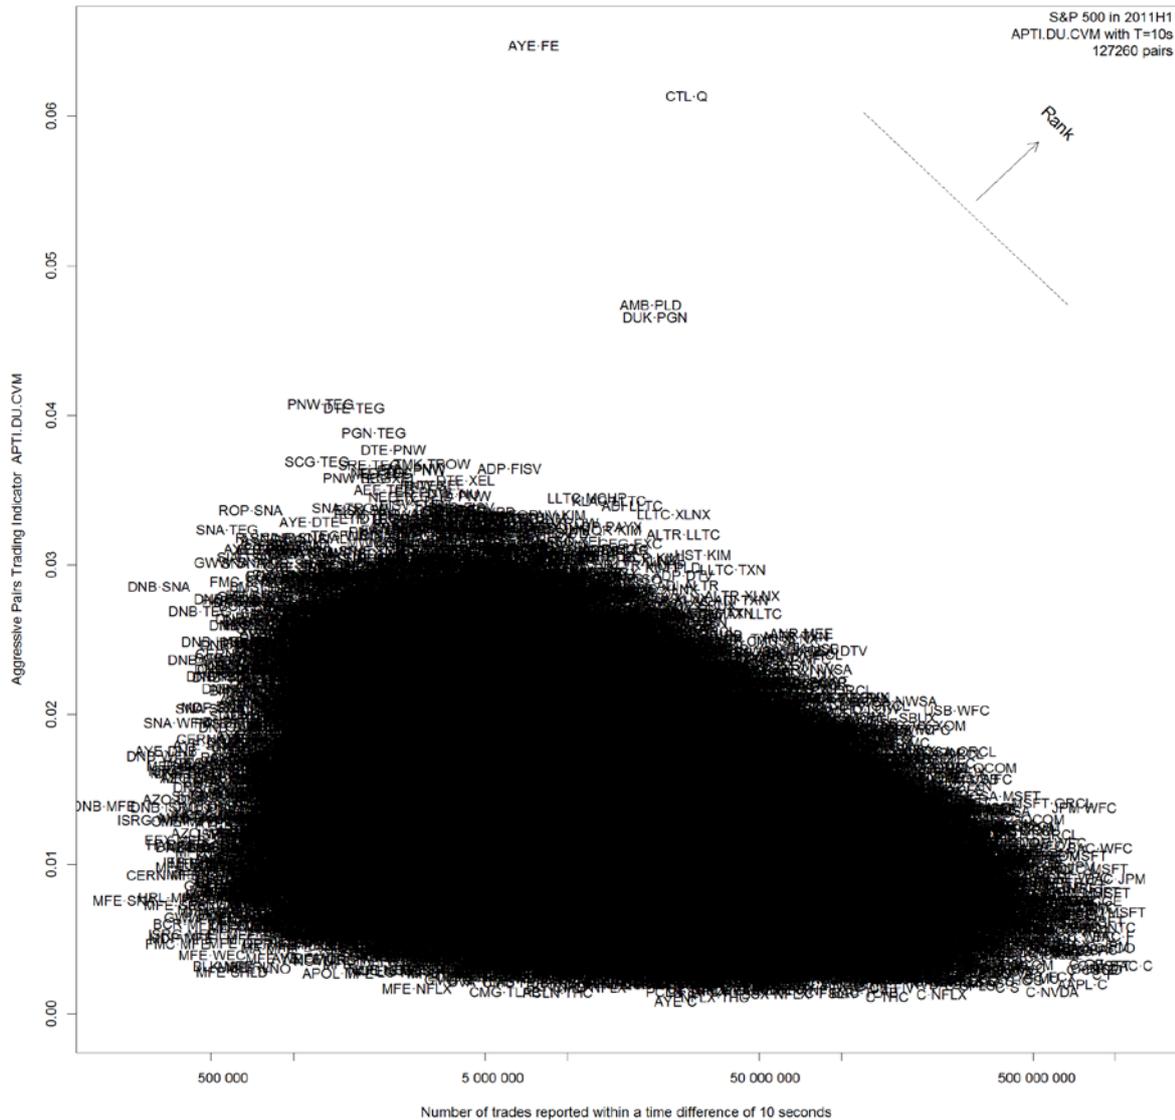
Figure 8 Scatterplot of aggressive pairs trading inferred empirically by indicator $APTI_{CVM}$ in 127 260 pairs of S&P 500 constituent combinations in the first six months of 2011. Pairs are formed between all combinations including index additions and removals during the period, provided the securities trade at least 1% as frequently as the most frequently traded security. When plotted in two dimensions with inferred related trading vertically and the frequency of trading horizontally, a frontier of pairs with the greatest inferred related trading faces the upper right corner. Examples of outstanding security pairs near the top of the diagram include AYE-FE and CTL-Q which were candidates for merger arbitrage trading during this period. Pairs along the leading (upper right) frontier tend to comprise related industries such as ALTR-LLTC (semiconductors), MSFT-ORCL (software), and JPM-WFC (banking). Detailed analysis of these pairs is listed in Table 3.

Table 3 shows an investigation of the pairs on the diagonal frontier. It finds the top three pairs identified (AYE-FE, CTL-Q, and AMB-PLD) are candidates for merger arbitrage because they merged during the study period. The next 25 listed along the diagonal frontier (such as PGN-TEG and LLTC-XLNX) are all strong candidates for pairs trading because their securities each turn out to be from the same industries. By contrast an inspection of the bottom end finds pairs such as C-NFLX (Citigroup and Netflix: a bank and a movie subscription service) that seem to have no apparent relation. The way the RTI and APTI approaches are finding strong candidates for pairs trading among the 127 260 pairs analyzed suggests they are successful at detecting such trading empirically. The methodology appears to be accurate at its task.

Pair	Potential explanation for the empirically inferred trading
AYE-FE	FirstEnergy Corp (FE) is an electricity provider which absorbed Allegheny Energy Inc (AYE) on 25 February 2011, two months into the period analyzed. The inferred trading is likely to be merger arbitrage.
CTL-Q	CenturyLink Inc (CTL) is an integrated telecommunication company which absorbed Qwest Communications International Inc (Q) on 1 April 2011 in the middle of the period analyzed. The inferred trading is likely to be merger arbitrage.
AMB-PLD	AMB Corporation (AMB) and Prologis Inc (PGD) were real estate investment trusts (REITs) that merged on 3 June 2011 toward the end of the analysis period. The inferred trading is likely to be merger arbitrage.
DUK-PGN PNW-TEG DTE-TEG DTE-PNW SCG-TEG DTE-XEL	Duke Energy Corporation (DUK), Progress Energy Inc (PGN), Pinnacle West Capital Corporation (PNW), Integrys Energy Group Inc (TEG), DTE Energy Co (DTE), SCANA Corporation (SCG) and Xcel Energy Inc (XEL) are energy companies with electricity generation, transmission, and/or distribution businesses. Any pair of these is a plausible candidate for pairs trading based on common industry fundamentals.
ADP-FISV	Automatic Data Processing (ADP) and Fisserv Inc (FISV) are providers of human payroll systems and financial services technology. Pairs trading is plausible based common industry fundamentals.
LLTC-MCHP KLAC-LLTC ADI-LLTC LLTC-XLNX ALTR-LLTC LLTC-TXN ALTR-XLNX	Linear Technology Corporation (LLTC), Microchip Technology Inc (MCHP), KLA-Tencor Corporation (KLAC), Analog Devices Inc (ADI), Xilinx Inc (XLNX), Altera Corporation (ALTR) and Texas Instruments Inc (TXN) are all high-tech companies involved in the integrated circuit industry. Any pair of these is a plausible candidate for pairs trading based on common industry fundamentals.
HST-KIM	Host Hotels & Resorts (HST) and Kimco Realty Corp (KIM) are REITs with interests in hotels and shopping centers respectively. Pairs trading is plausible based common industry fundamentals.
BK-USB USB-WFC BK-WFC JPM-WFC BAC-WFC BAC-JPM	The Bank of New York Mellon Corporation (BK), U.S. Bancorp (USB), Wells Fargo & Company (WFC), JPMorgan Chase & Co (JPM) and Bank of America Corporation (BAC) are all financial sector businesses providing banking, insurance, investments, and finance. Any pair of these is a plausible candidate for pairs trading based on common industry fundamentals.
CVX-XOM COP-XOM	Chevron (CVX), Exxon Mobil (XOM) and ConocoPhillips (COP) are oil and gas companies. Any pair of these is a plausible candidate for pairs trading based on common industry fundamentals.
CMCSA-NWSA	Comcast Corporation and News Corporation are providers of entertainment, information and news. Pairs trading is plausible based on the common industry.
MSFT-ORCL	Microsoft (MSFT) and Oracle (ORCL) are two of the world's leading software development companies. Pairs trading is plausible based on the common industry.

Table 3 Analysis and explanations for pairs trading in leading security pairs from Figure 8. The three pairs most outstanding (AYE-FE, CTL-Q, and AMB-PLD) are found to have merged during the period of analysis so trading inferred empirically is likely to have arisen from merger arbitrage. The next 25 pairs identified on the frontier are found to comprise companies in the same industries making them plausible candidates for pairs trading based on common fundamentals. The 28 leading pairs happen to be either involved in mergers or share industry fundamentals, despite comprising just 0.022% of the 127 260 pairs plotted in Figure 5. The fact the leading pairs identified from empirical inference of pairs trading have strong justification for merger arbitrage or pairs trading gives confidence the methodology is working.

The other three indices provide similar kinds of results. Investigation of the NASDAQ 100 constituent pairs in 2010 identifies LLTC-XLNX and MSFT-ORCL which are common with the S&P 500 analysis but offer robustness because the analyses here are for disjoint periods. Several company pairs are identified which turn out to be in the same industries. Investigation of the S&P MidCap 400 (which by construction, unlike the NASDAQ 100, has no overlap with the S&P 500) highlights pairs which turn out to be from the same industry and are strong candidates for pairs trading. Analysis of the FTSE 100 finds two outstanding pairs which each comprise two share classes from the same company (NG-NGn and SDR-SDRt) in addition to several pairs

found to be in the same industries. The methodology appears to be working as anticipated and is robust to the different exchanges and calendar periods involved. Further investigations suggest there is more to learn about applying the methodology to other exchanges internationally. Initial explorations of the German DAX 30 constituents and Australian ASX 200 constituents found little inference of pairs trading, compared with the US and UK index constituents. It would be interesting to test the reasons. Perhaps the DAX 30 contains too few constituents and perhaps pairs trading in the Australian market is less well developed than in the US and UK, or alternatively perhaps there is something different about the way those exchanges report their trades. There are many opportunities for subsequent research. It is sufficient here to demonstrate the potential of the methodology by discussing its success with constituent pairs from the S&P 500, S&P Midcap 400, NASDAQ 100, and FTSE 100 indices.

4. Applications of Pairs Trading Detection and Opportunities for Research

There are many avenues for research that become available given a method of detecting pairs trading, or more generally any related trading between two securities. I will outline four broad areas which arise from adjusting the partitioning of trades, the types of securities studied, the exchanges studied, and the frequency of computation.

4.1 Inferring dominant trading strategies

Pairs trading is a mean-reverting strategy, meaning it relies on the premise that a price difference or relative value will return to some kind of mean in the near future. A typical trading rule is to enter a position when the price difference is unusually high or low relative to a moving average, and to exit when it either returns to become close enough to that target or diverges to a stop loss. There are plenty of simulation studies, but these do not necessarily correlate with the trading actually taking place. The insight is to recognize that if a particular strategy is being traded, there should be an excess of entries detected when the entry signal reaches appropriate levels.

Time difference analysis and the related trading methodology can test for the entry and exit times by partitioning the trading data according to the levels of the entry signal of each strategy. An excess of trading in each direction should be inferred in the appropriate partitions. Ideally the strengths of the CDF deviations observed should coincide with the partitions of the trading strategy's entry signal if such strategy is being traded.

Initial tests on the twin ADRs of BHP Billiton suggest the dominant trading strategies exploit price excursions from a moving average of about 40 minutes. They strongly reject the presence of a strategy that enforces price equality, which makes sense as there is an ongoing price difference. On the other hand, initial studies of the Royal Dutch Shell ADRs infer strong enforcement of price equality, again consistent with their pricing. It is encouraging if the empirical approaches developed here can help confirm our theoretical expectations in situations like these.

4.2 Investigating the effects of company news releases on merger arbitrage trading

A second avenue for exploration is to study pairs trading over time by shortening the time windows in which it is aggregated. Although reducing the study window increases noise, the related trading approach seems to be capable of making inferences in study windows of a day.

This opens the possibility of testing the impact of company news releases on merger arbitrage, a form of pairs trading that involves a short position the acquirer and long position in the target. In merger arbitrage, the position can be held through to completion of a successful merger, or can be closed out earlier to take a profit or stop further loss.

Figure 9 shows the inference of aggressive pairs trading each day during the acquisition of Sun Microsystems (NASDAQ: JAVA) by Oracle Corporation (NASDAQ: ORCL). Disregarding the noise in the daily RTI signals, there is strong graphical evidence of increased pairs trading activity from the date of the merger announcement through to completion of the merger. This is consistent with expectations. Further analysis could indicate whether these trades are entries or exits of the merger arbitrage positions, or could look at cases where activity is detected ahead of the public announcements. The purpose of Figure 9 is to demonstrate the possibilities for further research in this area.
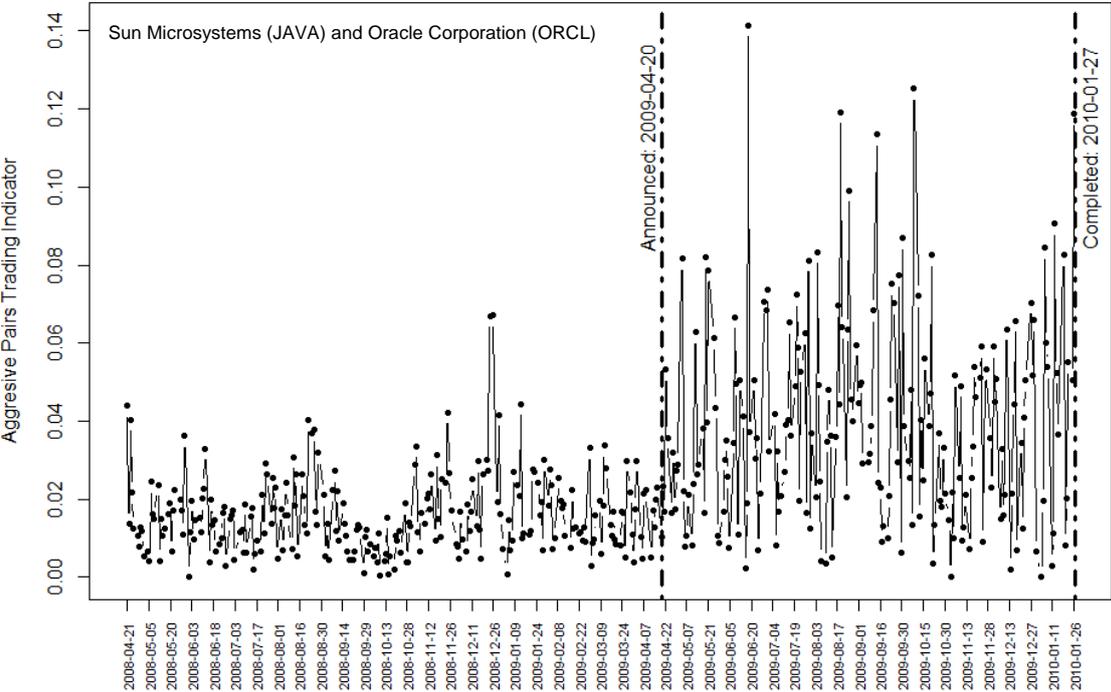

Figure 9 Daily inferences of pairs trading during the acquisition of Sun Microsystems (JAVA) by Oracle Corporation (ORCL). The merger was announced on 20 April 2009 and completed on 27 January 2010. Although computation of related trading indicators (RTIs) at daily frequency can be noisy as seen here, the methodology nonetheless infers a distinct increase in related trading activity from the date of the announcement through to completion of the merger. Further analysis could study whether the dominant type of merger arbitrage on each day is an entry or exit. This figure demonstrates the possibilities for further research into merger forensics with these techniques.

4.3 Studying timing differences between security exchanges

Financial securities can often be traded on several exchanges besides their primary listing. Differences in reporting mechanisms between these exchanges have the potential to affect the accuracy and precision of time difference analysis by introducing artifacts into the time differences observed. It may be possible to infer and eliminate these errors by studying asymmetries in the time differences observed. For example if the distribution function of BBL trading on exchange X followed by BHP trading on exchange Y changes shape when the

exchanges are reversed, ceteris paribus, a likely explanation could be different reporting times between the exchanges.

The robustness testing in section 3.1 which studied the five exchanges on which BHP and BBL were most frequently traded also led to a unique and repeatable ranking of the exchanges in terms of reporting speed. Inferring these reporting delays can open opportunities for research into exchange characteristics, as well as enabling the inferred delays to be removed from observed time differences to improve the precision of time difference analysis generally.

4.4 Broadening the security base to study rollover of derivatives

A fourth area for study is to expand the type of financial securities being considered. In this article I have looked only at stocks. Related trading can occur and be inferred in many classes of security including stocks, options, futures, bonds, commodities, and currencies. Related trading can also within each class and between classes, for example between a stock and its option, or a between a futures position and an underlying bond.

In financial derivatives it could be interesting to study the characteristics of futures rollovers each month where a maturing position is closed at the same time a position with longer maturity is opened. For example it may be possible to infer a relation between the "moneyness" of a position at rollover (how far it is into or out of the money) and the number of days remaining to maturity. Extending the rollover idea into options would add a dimension of strike price as it may be possible to identify conditions when options are rolled over to higher, lower, or similar strikes.

5. Conclusions

This paper helps overcome one of the great challenges in empirical financial trading research – the anonymity of trades – to infer related trading in two securities. Empirical inference of related trading enables insights into the nature of pairs trading in practice. The approach works by "time difference analysis" rather than "time series analysis", and more specifically by analyzing the statistics of time differences between trades. If the trading in two securities is unrelated, the empirical distribution of time differences should be approximately uniform, though in practice it is rare to see a uniform distribution empirically because there is so much other related trading taking place such as program and index trades.

Instead of using formal statistic tests which would give false positives in those circumstances, I use numerical related trading indicators (RTIs) which infer the relative amounts of related trading occurring. These indicators are a difference measure between the empirical distribution of trade time differences observed and the uniform distribution or null hypothesis proxy for unrelated trading.

The RTIs are found to be capable of distinguishing security pairs expected to be pairs traded from large pseudo-random sets of similar pairs. Aggressive pairs trading indicators (APTIs) incorporate buy and sell inferences from price information and can find securities pairs involved in merger arbitrage and pairs trading from amongst massively large sets of pairs of index constituents. Passive pairs trading indicators (PPTIs) and basket trading indicators (BTIs) are defined similarly to infer those kinds of trades.

The methodology is shown to be robust to calendar periods, different reporting delays between exchanges, intraday seasonal or diurnal variations, and the amount of background noise from program or index trading.

Applications include searching for actively traded pairs, inferring dominant entry and exit rules, and studying merger arbitrage. The ability to detect related trading in general may open more ways to study trading in two or more securities, for example to detect insider trading ahead of merger announcements, or to investigate the characteristics of futures and options rollover decisions. The literature on pairs trading is still in its infancy. Research opportunities abound, and much work remains to be done.

References

Alexander, C. and Dimitriu, A. (2002) The cointegration alpha: Enhanced index tracking and long-short equity market neutral strategies, ISMA Discussion Papers in Finance 2002-08, ISMA Centre, University of Reading, UK.

Altucher, J. (2004) *Trade Like a Hedge Fund*, John Wiley & Sons Inc., Hoboken, New Jersey.

Bookstaber, R. (2007) *A Demon of our Own Design*, John Wiley and Sons, Hoboken, New Jersey.

Bowen, D. A., Hutchinson, M. C. and O'Sullivan, N. (2010) High-frequency equity pairs trading: transaction costs, speed of execution, and patterns in returns, *The Journal of Trading*, 5(3): 31-38.

Chen, H., S. J. Chen and Li, F. (2010) Empirical investigation of an equity pairs trading strategy, *SSRN working paper 1361293*, April 2010.

Deardorff, A. V. (1979) One-way arbitrage and its implications for the foreign exchange markets, *Journal of Political Economy*, 87(2).

Do, B. and Faff, R. (2009) Does simple pairs trading still work? *Working paper*, 10 Feb 2009.

Dunis, C. L., Giorgioni, G, Laws, J. and Rudy, J. (2010) Statistical arbitrage and high-frequency data with an application to Eurostoxx 50 equities, *working paper*, March 2010.

Engle, R.F., and Russell, J.R. (1998) Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data. *Econometrica* 66(5), 1127-1162.

Froot, K.A. and Dabora, E.M. (1999) How are stock prices affected by the location of trade? *Journal of Financial Economics* 53(2), 189-216.

Gatev, E., Goetzmann, W.N., and Rouwenhorst, K.G. (2006) Pairs trading: Performance of a relative-value arbitrage rule, *Review of Financial Studies* 19(3), 797-827.

Gouriéroux, C., Jasiak, J. and Le Fol, G. (1999) Intra-day market activity. *Journal of Financial Markets* 2(3), 193-226.

Hedge Fund Research (2011) *HFR global hedge fund industry report: Second quarter 2011*, retrieved 29th February 2012 from http://www.hedgefundresearch.com/pdf/hfr_ir_sample.pdf

Hong, G. and Susmel, R. (2003) Pairs-trading in the Asian ADR market, unpublished working paper, Saginaw Valley State University, and Bauer College of Business, University of Houston.

Jacobs, B. I. and Levy, K. (2005) *Market Neutral Strategies*, Hoboken, New Jersey, John Wiley & Sons.

Lee, C.M.C. and Ready, M.J. (1991) Inferring trade direction from intraday data, *Journal of Finance* 46(2), 733-746.

Livermore, J. L. (1940) *How to Trade in Stocks*, Duell, Sloan, and Pearce, New York.

Nath, P. (2003) High frequency pairs trading with U.S. treasury securities: risks and rewards for hedge funds, *SSRN working paper 565441*.

Pacurar, M. (2006) Autoregressive Conditional Duration (ACD) Models in Finance: A Survey of the Theoretical and Empicial Literature, *SSRN working paper 933120.*

Paul, S. (2008) Australia eases shorting ban for dual-listed stocks, *Thomson Financial News Limited*, 23rd September 2008.

Reverre, S. (2001) The Complete Arbitrage Deskbook, McGraw-Hill, New York, pp 474-480.

Rosenthal, L. and Young, C. (1990) The seemingly anomalous price behavior of Royal Dutch / Shell and Unilever N.V./PLC, *Journal of Financial Economics* 26(1), 123-141.

Schultz, P. and Shive, S. (2010) Mispricing of dual-class shares: Profit opportunities, arbitrage, and trading, *Journal of Financial Economics*, 98(3): 524-549.

Suarez, E. D. (2005) Arbitrage opportunities in the depositary receipts market: Myth or reality? *Journal of International Financial Markets, Institutions and Money* 15: 469-480.

Wojcik, R. (2005) *Pairs trading: a professional approach*, PDF dated 19 November 2005, retrieved in 2012 from http://www.futuresindustry.org/downloads/Audio/Companion/Three-2812.pdf