To lead or to lag? Measuring asynchronicity in financial time-series using dynamic time warping*

Clint Howard[1,2], Tālis J. Putniņš[1,3] and Vitali Alexeev[1,4]

[1] *University of Technology Sydney*
[2] *Robeco Institutional Asset Management*
[3] *Stockholm School of Economics in Riga*
[4] *University of Guelph, Canada*

September 10, 2022

Abstract

Financial time-series are rarely perfectly synchronized. More commonly, one market or one security leads another, causing mis-estimation in empirical modeling. A particularly challenging aspect of the asynchronicity is that the lead/lag changes dynamically through time and with market conditions. We show how dynamic time warping (DTW) can be used to measure dynamic asynchronicity and improve empirical models in asset pricing and price discovery. Simulations confirm the validity of the DTW approach for a range of scenarios. We show that covariance measures estimated using DTW to correct dynamic asynchronicity recover a positive relation between market risk and return, helping resolve the beta anomaly. At intraday horizons, applying DTW shows that global market price discovery oscillates between S&P 500 futures and FTSE 100 futures, whereas standard measures miss this important intraday dynamic.

JEL classification: G12, G14, G17

Keywords: dynamic time warping, asset pricing, beta anomaly, price discovery, asynchronous data

*Corresponding author: Clint Howard. E-mail address: clint.s.howard@student.uts.edu.au
We thank David Blitz, Harald Lohre, and seminar participants at Robeco Institutional Asset Management for helpful comments and suggestions. The author[2] is employed by Robeco Institutional Asset Management. The views presented here are their own and do not necessarily reflect the views of their employer. Computational facilities were provided by the UTS eResearch High Performance Computer Cluster.

1. Introduction

Asynchronicity in financial time-series poses a significant challenge in financial markets research. Empirical models commonly assume that all observations occur contemporaneously or at a fixed time lag. Lo and MacKinlay (1990) demonstrate that frictions in markets can impede information transmission, leading to lead/lag patterns between observable asset time-series data. Frictions due to infrequent trading, information flowing from large stocks to smaller stocks, and even the physical distance between trading venues can all contribute to asynchronicity between observed time-series data. This asynchronicity between asset returns can cause errors in inference. For example, it can result in mis-estimation in price discovery models, and measurements of asset covariance can be biased due to asynchronicity in the underlying time-series. The dynamic nature of asynchronicity presents the biggest challenge where extant econometric frameworks fail. We use dynamic time warping (DTW) to correct for this dynamic asynchronicity between financial time-series.

DTW is an alignment algorithm which provides a measure of similarity between two time-series, first utilized in speech recognition (Sakoe and Chiba, 1978; Keogh and Pazzani, 2002). The power of DTW is that it allows for a flexible alignment of both leading and lagging patterns between two or more time-series. This feature of DTW is well-suited to our problem of measuring asynchronicity between financial time-series and accounting for the time-varying nature of the asynchronicity. With DTW we do not need to specify any structural assumptions on the expected behavior of the lead/lag, the algorithm itself is able to freely uncover dynamic lead/lag structures between time-series.

We validate the use of DTW to recover simulated lead/lag structures across several simulated time-series scenarios. Across these simulated lead/lag scenarios, DTW is able to successfully recover the induced lead/lag patterns with an average mean absolute error ($\widetilde{MAE}$) of 5.9%. We find that the success of DTW in recovering the lead/lag pattern is primarily a function of noise contaminating the time-series, and that the success is stable for varying levels of volatility of the simulated time-series. This is an important result, suggesting that DTW can be used for assets that are inherently volatile. Whilst in our simulated scenarios we are able to control the noise, volatility and lead/lag between time-series, the true test for the DTW method is applying it to varying empirical scenarios.

Real-world asynchronicity in financial time-series can occur at many different frequencies, from the macro to the micro. To demonstrate DTW's relevance for financial time-series, we select two empirical applications. One at the daily frequency, and one at the intraday frequency. Our first application is in the measurement of a stock's beta to the market. Under the Capital Asset Pricing Model (CAPM) contemporaneous stock and market returns are used to estimate a stock's beta. The CAPM then predicts that there is a positive relation between a stock's beta and expected returns. Several studies (Reinganum, 1981; Lakonishok and Shapiro, 1986; Fama and French, 1992) present empirical evidence that there is, in fact, a negative relation between risk and return, giving rise to the beta anomaly. The literature has typically approached this empirical anomaly in two ways. The first approach focuses on the mis-estimation of beta arising from the regression framework used (Dimson, 1979; Scholes and Williams, 1977). The second approach focuses on anomaly-based explanations such as betting against beta (BAB) and the low-volatility effect (Frazzini and Pedersen, 2014; Blitz and van Vliet, 2007; Blitz et al., 2019)), lottery premia effects(Bali, Brown, Murray and Tang, 2017) and idiosyncratic volatility (Liu, Stambaugh and Yuan, 2018).

We use DTW to extend the work of Dimson (1979) and Scholes and Williams (1977) by incorporating a dynamic asynchronicity adjustment in the estimation of beta, without imposing a strict assumption on the lead/lag structure. Both the Dimson (1979) and Scholes and Williams (1977) beta estimations use various leading/lagging market return variables as additional independent variables in the beta estimation regression. However, both methods assume that the asynchronicity between the market return and stock return is static within the defined leading/lagged variables. The dynamic of the lead/lag, however, need not necessarily be static and is potentially time-varying. By using DTW to better align stock returns and market returns, we allow for a more flexible incorporation of asynchronicity into the estimation of beta.

We find that the mis-estimation of beta is primarily driven by small stocks, giving further strength to the non-synchronous trading argument for the mis-estimation of beta. Smaller, less liquid stocks take longer to incorporate market-wide information and thus lag behind the market. By using DTW to account for the dynamic asynchronicity between stock returns and market returns, we are able to more accurately estimate a stock's beta and help resolve the beta anomaly. However, we cannot claim full resolution of the beta anomaly. From 2000, we see that across all estimates of beta that high beta stocks underperform low beta stocks. This suggests that there are other factors contributing to the beta anomaly, in addition to a non-synchronous trading component.

Our second application applies DTW to study the intraday dynamics in price discovery between two instruments. Established measures of price discovery, such as Hasbrouck (1995)'s Information Share (IS) and Gonzalo and Granger (1995)'s Component Share (CS), typically provide a summary statistic of price leadership between two instruments across an estimation window. These methods do not provide insight into the dynamics of the price discovery process within the estimation window. Ozturk, van der Wel and van Dijk (2017) propose a novel approach using flexible Fourier transformations to measure the the intraday dynamics in IS by allowing for time-varying volatility of the efficient price innovations and idiosyncratic noise. By applying DTW we obtain an estimate of the lead/lag at each time-step in the estimation window, allowing a clearer insight into the intraday temporal dynamics between time-series.

We explore intraday price discovery dynamics across global index futures. Using DTW we demonstrate rich temporal behavior of the intraday lead/lag that is anchored around significant market operation events, such as the market open and close in the underlying equity markets. We select the E-mini S&P 500 futures and FTSE 100 futures, respectively traded on the Chicago Mercentile Exchange (CME) and London Stock Exchange (LSE), as these contracts represent two of the largest and most liquid equity markets. Our results show that the lead/lag structure between the two index futures has evolved over time. At the daily level we document a compression in the average lead/lag towards zero from 2001–2010. As markets have become increasingly automated the lead/lag between the two contracts is more likely to be restricted by the speed of information transmission.

Through our two empirical applications, we show that asynchronicity is highly dynamic and can have a material effect on inference in empirical models of asset pricing, risk, and price discovery. Using DTW to account for dynamic asynchronicity, we generate new insights and resolve existing empirical challenges in the literature. Our approach provides a pathway to the further study of existing and new problems where the asynchronicity inherent in financial time-series obfuscate model estimation and inference. One such problem is in the emerging literature around the recent increase in co-movement in markets (e.g., from high-frequency trading (Malceniece, Malcenieks and Putniņš, 2019)). It is currently difficult to attribute this increase in co-movement to a genuine change in the systematic risk in markets. An alternative explanation proposes that it is due to a better alignment of returns in assets within markets that arises as a result of the increased liquidity and trading activity afforded to stocks when they're added to an index (Barberis, Shleifer and Wurgler, 2005). Li, Yin and Zhao (2020) explore the effects of program trading on the co-movement of stocks, and document evidence that stocks preferred by algorithmic traders exhibit excessive co-movement above what would be expected by fundamental drivers of return. This is an example where DTW could be used to further study the drivers of the empirical results which were documented, and potentially provide a richer insight into the nature of the effect of program trading on co-movement by removing the dynamic asynchronicity that is present in returns.

The finance literature on asynchronicity in time-series has evolved as the speed of information transmission has increased. Earlier studies (Scholes and Williams, 1977; Dimson, 1979; Kawaller et al., 1987) typically use lagged variables in a regression framework to account for the intertemporal relation between financial time-series. As trading in markets has become predominantly automated, and the speed of trading has increased, the use of regression frameworks to estimate lead/lag is often inadequate as the lead/lag relationship has reduced from minutes, to seconds, to milliseconds. In the past 20 years, new techniques have emerged for measuring high-frequency lead/lag. Hayashi and Yoshida (2005) develop a covariance estimator for non-synchronous diffusion processes which is used to account for temporal dynamics between time-series. Dobrev and Schaumburg (2016) present a model-free method for estimating the lead/lag relationship using a timing offset between trading activity to estimate which market is driving the lead/lag. The primary difference with our approach using DTW, is that we are able to obtain point-in-time estimates of the lead/lag relationship at every observation within the estimation window. This allows us to examine the time-varying dynamics in the lead/lag relationship across the estimation window, which is often not possible with the covariance estimator of Hayashi and Yoshida (2005).

Our work also contributes to the strand of literature on synchronizing returns across markets with respect to different market closing times. When modeling assets that are traded in different markets you must account for the difference in market closing times, otherwise any estimates which are derived from time-series of these assets could result in significant mis-estimation of models. A significant body of literature propose various methods which aim to synchronize returns across different markets using a variety of statistical models (Burns, Engle and Mezrich, 1998; Martens and Poon, 2001; Audrino and Bühlmann, 2004; Scherer, 2013). DTW is a direct complement to these approaches, and could be applied in the synchronization of stock returns across different markets.

Ito and Sakemoto (2020) and Franses and Wiemann (2020) are perhaps the most closely related papers to ours, using DTW to study lead/lag relationships in foreign exchange markets and business cycles, respectively. Ito and Sakemoto (2020) propose using DTW to measure high-frequency foreign exchange lead/lag patterns. Similar to our work exploring the dynamics of intraday price discovery, they document a clear change in the lead/lag patterns between currency pairs in response to important market announcements. We expand on their work by further exploring the sensitivity of the optimal series alignment based on DTW to noise and volatility, and provide further empirical evidence on the use of DTW for studying asynchronicity in financial time-series.

Our paper provides a new method for estimating lead/lag between financial time-series, and a mechanism for incorporating this estimation of lead/lag to synchronize time-series. Using this method we demonstrate that the beta anomaly can be resolved when properly accounting for asynchronicity between stock returns and market returns. Our result brings into question some of the recent literature on the BAB phenomenon Frazzini and Pedersen (2014) and is more closely aligned with the work of Novy-Marx and Velikov (2022) which call into question some of the exceedingly strong performance of the BAB factor. At the intraday frequency we show clear patterns of intraday lead/lag effects between E-mini and FTSE 100 futures, highlighting rich temporal dynamics that traditional price discovery measures are unable to capture. This provides evidence that better measures of price discovery, which can appropriately incorporate changes in price leadership within the estimation window, are needed.

2. Validation of DTW approach

2.1. Simulation design

We use simulation to validate and study the effectiveness of our proposed method for measuring the lead/lag between financial time-series. When applying DTW, it is important that the time-series share some commonality, as the DTW algorithm will find a relationship where there is none. To validate DTW, we follow Putniņš (2013) and use a time-series model where two stocks share a common fundamental value. The fundamental value is assumed to follow a random walk:

$$m_t = m_{t-1} + u_t,\ u_t \sim N(0, \sigma_u), \tag{1}$$

where m_t is the natural logarithm of the fundamental value at time t, and u_t is a noise component. We define a time-series p_i, which tracks the fundamental value with a state-dependent time-shift of $\delta_{i,t}$ periods, and noise, $s_{i,t}$, as:

$$p_{i,t} = m_{t-\delta_{i,t}} + s_{i,t},\ s_{i,t} \sim N(0, \sigma_{s_i}). \tag{2}$$

To test the efficacy of the DTW algorithm at capturing time-varying asynchronicity, we use three lead/lag scenarios which are representative of common lead/lag behavior in financial markets: constant, gradual, oscillating, and, for robustness, a randomly switching lead/lag. The constant lead/lag scenario represents a pair of stocks that have a constant structural lag between them, often arising from trading on multiple venues where one venue dominates the price discovery process. The gradual lead/lag represents a pair of stocks where, over the course of the trading day, the lag between the two stocks decreases. This may occur due to overnight information imbalances which shrink as information is incorporated across the trading day. The oscillating lead/lag represents a pair of stocks where, over the course of the trading day, the lag alternates between positive and negative. Such a situation may arise from the opening of other international markets and the impounding of new information across markets. The randomly switching lead/lag represents a pair of stocks where there is no predetermined dynamic asynchronicity, and thus the lead/lag relation between the assets may switch randomly over the trading day. This scenario is designed to test the robustness of the DTW algorithm at capturing any dynamic asynchronicity pattern.

To implement these four scenarios, we set p_1 as the reference time-series and p_2 as the secondary time-series. We vary $\delta_{2,t}$ to induce the desired lead/lag between the two time-series across the simulation. With $\delta_{1,t}, \delta_{2,t} \in \mathbb{N}$, we have:

$$p_{1,t} = m_{t-\delta_{1,t}} + s_{1,t}, \tag{3}$$

$$p_{2,t} = m_{t-\delta_{2,t}} + s_{2,t}. \tag{4}$$

A constant lead/lag is induced by fixing δ_2 to a constant value across the simulation. This produces a true lag of $\delta_2 - \delta_1$. A gradual lead/lag is induced by setting a maximum value for $\delta_{2,t} = \delta_2$ and then creating a sequence of evenly spaced lags $\delta_{2,t} = \left\lfloor \delta_t - \frac{\delta_2(N-t)}{N} \right\rceil$ for $t = 0, \ldots, N$ where N is of cardinality $\{p_{i,t}\}$. An oscillating lead/lag takes the functional form $\delta_{2,t} = \lfloor A \sin(2\pi x_t) \rceil$ where $A \in \mathbb{N}$ is a scaling factor, and x_t is taken from an evenly spaced grid of values, $x = \{x_0, x_1, ..., x_t\} = \left\{0, \frac{1}{N}, ..., 1\right\}$. This induces a sinusoidal wave of period 2π and amplitude of one. We scale by A to ensure the $\delta_{2,t}$ are proportional to δ_1. A randomly switching lead/lag is induced by taking both $\delta_{1,t}$ and $\delta_{2,t}$ for each time-step t, and for each $\delta_{i,t}$ we randomly sample a variable $z \sim N(0, 1)$ and if $z > 3.7$, we draw $\delta_{i,t}$ from $\mathbb{N}(20)$ and set this as the new δ_i for p_i.

2.2. Dynamic time warping

We present the generalized version of DTW where the two time-series can take differing lengths, N and M, such that $p_{1,t} = (p_{1,1}, p_{1,2}, ..., p_{1,n})$ and $p_{2,t} = (p_{2,1}, p_{2,2}, ..., p_{2,m})$. We compute an $N \times M$ cost matrix, $C \in R^{N \times M}$, where each element in row n and column m corresponds to a distance between the elements $(p_{1,n}, p_{2,m})$. Typical distance measures used include the Euclidean distance and Manhattan distance, we use the Euclidean distance.

The objective of the DTW algorithm is to find the optimal alignment between p_1 and p_2 which has the minimum overall cost. We define a warping path $Z = (z_1, z_2, ..., z_K)$ with $z_k = (n_k, m_k)$ as the set of matrix elements that define a mapping between p_1 and p_2, in which the following conditions are satisfied:

- Boundary condition: $z_1 = (1, 1)$ and $z_k = (n, m)$

- Monotonicity condition: $n_1 \leq n_2 \leq \cdots \leq n_K$ and $m_1 \leq m_2 \leq \cdots \leq m_K$

- Step-size condition: $z_{l+1} - z_l \in \{(1,1),(1,0),(0,1)\}$

The optimal warping path is the path that has the minimum total cost among all possible warping paths. Dynamic programming methods are used to calculate an accumulated cost matrix, D. Each element in the accumulated cost matrix is defined as the local cost measure in the current cell, $c(x_n, y_m)$, plus the minimum of the accumulated cost measure in adjacent cells:

$$D(n,m) = \min(D(n-1,m-1),\ D(n-1,m),\ D(n,m-1)) + c(x_n, y_m). \tag{5}$$

Using the accumulated cost matrix, the optimal warping path Z^* is calculated by stepping in reverse index order through D using the following algorithm:

$$w_{k-1} = \begin{cases} (1,\ m-1) & \text{if } n = 1, \\ (n-1,\ 1) & \text{if } m = 1, \\ \arg\min(D(n-1,m-1),\ D(n-1,m),\ D(n,m-1)) & \text{otherwise.} \end{cases}$$

Given the empirical settings we intend to test, we impose and test a set of global constraints. We use the classic Sakoe-Chiba band which runs along the main diagonal and has a fixed width $W \in \mathbb{N}$. This is a global constraint, which implies that an element $p_{1,N}$ can only be aligned to some value $p_{2,N}$ where $m \in \left[\frac{m-W}{n-W}(N-W),\ \frac{m-W}{n-W}(N+W)\right] \cap [1:m]$. We also loosen the boundary condition and allow $z_1 \in (1,1) : (1+\psi,\ 1+\psi)$ and $z_k \in (n-\psi,\ m-\psi) : (n,\ m)$ where $\psi \in \mathbb{N}$. We loosen the boundary condition so that we do not require the start and end points of the time-series to align perfectly at the start and end of the trading day, as is often the case for high-frequency financial data.

In our setting, we use the optimal warping path Z^* as a measure of the asynchronicity between p_1 and p_2 at each time-step t. The optimal warping path will have at minimum a length of $\max(N, M)$, this can result in there being more lags in the optimal warp-path than time-steps. This occurs due to index duplication where the lag does not change. To map from warp-time back to clock-time we take all index pairs $z_k = (n_k, m_k)$ and for each duplicated value of m_k we take the first occurrence of m_k and map this to time-step k. The asynchronicity at each time-step is then calculated as:

$$\delta_{k_{DTW}} = n_k - m_k. \tag{6}$$

We define a measure of the algorithms performance as how well it recovers the true lag, δ_{True}, of the simulated time-series by using the mean absolute error scaled by the average true lag across the simulation:

$$\widetilde{MAE} = \frac{\delta_{k_{DTW}} - \delta_{k_{True}}}{\frac{1}{N}\sum_{i=1}^{N}\delta_{i_{True}}} \quad if\ \frac{1}{N}\sum_{i=1}^{N}\delta_{i_{True}} \neq 0 \tag{7}$$

2.3. Simulation results

Figure 2 shows the induced lead/lag structures between the two time-series and the lead/lag recovered from the DTW algorithm in the four scenarios. We find that in each of the four scenarios, the DTW algorithm is successful at recovering the simulated lead/lag pattern.

<Insert Figure 2 about here>

Table 1 presents several metrics from the four baseline simulations. Across all four cases we find consistent results. In the baseline scenarios we use a lag of negative ten ticks (i.e., the difference in speed between p_1 and p_2 is ten ticks) and using DTW we recover this induced lag with a $\widetilde{MAE}$ between 5.83% and 6.50%. The error in DTW's ability to recover the lag is a combination of bias in the underlying DTW algorithm and noise in the two time-series. The bias in the DTW algorithm is a function of how the algorithm stitches the two time-series together. DTW, like many algorithms, is prone to biases. The DTW algorithm tries to minimize the total cumulative distance between the two time-series, however this can result in localized stitching errors due to selection choices that the algorithm has to make[1]. By switching which time-series is the primary time-series in the DTW algorithm, we find

[1]Table A.1 in the Internet Appendix examines the characteristics of this bias

a directional bias. We correct for this in our subsequent analyses by running a two-step DTW, in which we alternate which time-series is the primary time-series, and then take the difference of the final results divided by two. From Table A.1 we can see that after adjusting for the bias, the DTW algorithm successfully recovers the induced true lead/lag with minimal error.

<Insert Table 1 about here>

In addition to standalone simulations, we use bootstrapping to study the effects of the simulation parameters on the accuracy of the DTW algorithm in recovering the induced lead/lag patterns. We run 10,000 simulations where we apply the DTW algorithm on two simulated time-series under different parameter combinations and record the $\widetilde{MAE}$. In Figure 3 we vary δ_2 and σ_2 which allows us to study the accuracy of DTW when we alter the absolute level of the lag and the ratio of noise between the two time-series. This is of relevance to applications where there might be a difference in the volatility of the assets being studied. We vary δ_2 between zero and ten, and σ_2 between zero and two for each of the four lag scenarios. In these simulations, we fix $\delta_1 = 5$, so when $\delta_2 = 5$, the induced lag is zero. We observe that when the induced lag is zero, $\widetilde{MAE}$ is highest. This is a result of the directional bias in the DTW algorithm, which produces a larger $\widetilde{MAE}$ when the true lag is close to zero. We also observe this effect in the gradual lag and switching lag scenarios, noting that in the gradual lag, this occurs when $\delta_2 = 6$ due to how the gradual lag structure is induced. In the oscillating scenario, we observe that the results are primarily a function of σ_2, as the $\widetilde{MAE}$ is relatively constant for different levels of δ_2. This is due to the symmetric nature of the bias in the DTW algorithms, the oscillating lag is constructed to have an area under of the curve of zero (i.e., the lag structure is symmetric), we observe that the bias cancels out, leaving an approximately linear relation between $\widetilde{MAE}$ and σ_2. Overall, we conclude that error in the DTW algorithm is primarily a function of the relative noise in the time-series, and that as the level of noise in one time-series increases relative to another time-series, the error will also increase. It is important to take this into consideration when applying DTW, but does not preclude the use of DTW when studying lead/lag in highly volatile financial time-series

<Insert Figure 3 about here>

Having established that the relative level of noise between the two time-series is a driver of error in the DTW algorithm, we explore the effect of three additional parameterizations on the accuracy of the DTW algorithm. Figure 4 presents three different parameter pair combinations for the gradually decreasing lag scenario. In panel A, we measure the $\widetilde{MAE}$ whilst varying σ_{s_1} and σ_{s_2}. We observe that the $\widetilde{MAE}$ is approximately linearly increasing in both σ_{s_1} and σ_{s_2}. This result is not unexpected, both σ_{s_1} and s_2 control the level of noise in the underlying time-series and the more noise that is induced, the more errors the DTW algorithm will make when it attempts to stitch the two time-series together. As we aim to apply DTW to measure lead/lag at high-frequency, the algorithm needs to be stable for time-series with large N. To test this, we measure the $\widetilde{MAE}$ for varying levels of N and σ_u. In panel B we observe that the $\widetilde{MAE}$ is decreasing in $\log(N)$ and constant in σ_u. $\widetilde{MAE}$ is decreasing in $\log(N)$, as N grows larger the relative effect of the errors is overall lower. $\widetilde{MAE}$ is constant in σ_u. σ_u is common to both time-series and thus the only variability which is induced is in the noise of the two time-series relative to the fundamental value, not in the fundamental value itself. In panel C, we observe a non-linear relation between σ_{s_2} and σ_u. We find that at lower levels of σ_{s_2}, the $\widetilde{MAE}$ is higher. This is likely due to the low signal to noise ratio which is created between the fundamental value and the time-series. This result suggests that if the fundamental value itself is highly volatile, the DTW algorithm will have a higher level of error. The results suggest that the DTW algorithm is more successful at recovering the lead/lag when the two time-series have a similar level of noise. This result has implications for the usage of DTW in price discovery applications, suggesting that the DTW algorithm would perform better when the volatility levels in the two time-series are similar.

<Insert Figure 4 about here>

3. A better beta

The CAPM predicts that stocks with high betas will outperform stocks with low betas (Black, Jensen and Scholes, 1972; Fama and MacBeth, 1973), however several empirical studies (Reinganum, 1981; Lakonishok and Shapiro, 1986; Fama and French, 1992) provide evidence that stocks with low betas outperform stocks with high betas. This is known as the "beta anomaly" or "low-beta effect". Being one of the key empirical challenges to the CAPM, there is a rich literature which attempts to explain the beta anomaly. There are typically two approaches. The first stream posits that the method used to estimate beta is not sufficient, and they present alternative construction methods, such as the approaches of Dimson (1979) and Scholes and Williams (1977). The second focuses on anomaly-based explanations, and pose that once controlling for other effects the beta anomaly is resolved. Liu et al. (2018) proposes a resolution for the beta anomaly by controlling for idiosyncratic volatility, whilst Bali et al. (2017) argue that investor demand for lottery-like stocks is a key driver of the beta anomaly. Whilst these two types of approaches control for and render the beta anomaly insignificant, they do not fully recover the expected relationship that the CAPM predicts, that stocks with high beta outperform stocks with low beta. We contribute to the literature by providing an alternative method to estimate beta, which takes into account dynamic asynchronicity between stock returns and the market returns. This approach recovers a weak positive risk-return relationship, which is robust to controlling for well-established asset pricing factors, including idiosyncratic volatility and the MAX effect.

<Insert Figure 1 about here>

A first question when applying DTW to estimating beta is whether we expect asynchronicity between stock returns and market returns. Figure 1 shows the proportion of the correlation between the market return and all stocks in the CRSP sample using one-day forward and one-day lagged stock returns, relative to the correlation using one-day forward, contemporaneous and one-day lagged stock returns. In effect, this is an approximate measure of the degree of asynchronicity in U.S. equity markets. Figure 1 demonstrates the dynamic nature of asynchronicity, as we see significant time-variation in the proportion of non-contemporaneous correlation. We observe that the time-varying nature of asynchronicity is not only a feature of small stocks, but is present across all capitalization levels. Empirical models have traditionally addressed this non-contemporaneous correlation by using leading/lagging indicators in a regression framework to account for asynchronicity. This approach, however, imposes a pre-specified and fixed structure around the nature of asynchronicity, and does not account for any time-dynamics in asynchronicity. Using DTW abolishes this confining assumption.

In addition to examining non-contemporaneous correlations, we can explicitly measure the lead/lag between stock returns and market returns using DTW. Table 2 presents the estimated lead/lag between stock returns and market returns for the standard CRSP sample. We find intuitive results, particularly for small stocks, where small stocks have gone from lagging the market by 0.81 days, on average, to approximately no lag alongside a reduction in the spread of lead/lag across small stocks. An important result is that the lead of large stocks has not significantly changed over our sample, which suggests that the asynchronicity that the DTW method is adjusting for tend to not be present in large stocks.

<Insert Table 2 about here>

One possible driver of the beta anomaly is that small, illiquid stocks are slower to incorporate market-wide information and using contemporaneous market returns to estimate beta results in a downwards biased beta. There are several methods which partially resolve this non-synchronous trading effect. Dimson (1979) and Scholes and Williams (1977) uses contemporaneous market returns combined with pre-specified forward and lagged market returns in the estimation of beta. Our approach uses DTW to align the stock return series with the market return series, taking into account any dynamic asynchronicity, and we then calculate the beta using the aligned returns series.

3.1. Data and method

Daily and monthly stock data used are from the Center for Research in Security Prices (CRSP). We take all U.S. stocks traded on all exchanges that are classified as ordinary common shares (SHRCD 10 or 11) with PRC greater than \$5, and adjust returns for delisting bias as per Shumway (1997) and Shumway and Warther (1999). Daily and monthly market returns, HML, SMB, RMW, CMA, UMD factor returns, and risk-free (one-month Treasury Bill) rates are from Kenneth French's data library. Daily FMAX factor returns are from Turan Bali's website (Bali et al., 2017). Balance sheet information which is used to calculate the book-to-market ratio is sourced from Compustat. Our sample covers the months t from July 1927 through November 2019, and we use month $t+1$ returns from August 1927 through December 2019.

Using DTW, we are able to systematically select the best alignment of stock returns and market returns for the regression, and remove the effect of asynchronicity on the ex-post estimation of beta. From our DTW approach, we calculate a series of stock characteristics, the raw DTW beta (β_{DTWR}), a bootstrapped DTW beta ($\hat{\beta}_{DTW}$), DTW beta (β_{DTW}), and the bootstrapped DTW t-statistic (β_{DTWT}). β_{DTWR} is estimated directly from the aligned market returns and stock returns. To address the inherent bias in the DTW method, as described in Section 2.3, we use bootstrapping. We permute the stock returns series, apply DTW to align the permuted stock returns and market returns, and then re-estimate $\hat{\beta}_{DTW}$. We then have the de-biased DTW beta which is the difference between the β_{DTWR} and $\hat{\beta}_{DTW}$. In addition to the DTW estimates of beta, we also calculate the bootstrapped t-statistic associated with estimating $\hat{\beta}_{DTW}$. The full calculation details can be found in Appendix A.

3.2. Results and discussion

<Insert Table 3 about here>

Table 3 presents the time-series average of the cross-sectional summary statistics of our DTW beta factors, the standard CAPM beta estimate (β), Dimson beta (β_{DIM}), Scholes-Williams beta (β_{SW}) and a set of control factors. Full details on factor construction can be found in Appendix A. β_{DTWR} has a mean greater than the three standard measures of beta, suggesting that there is a downwards bias in beta which can be addressed by adjusting for asynchronicity when estimating beta. One concern with this approach is the bias which exists in the DTW algorithm, and it's ability to find spurious correlation where it may not exist. Whilst β_{DTWR} may be a fair representation of beta, we find a strong relation with idiosyncratic volatility and that increases in the level of beta measured by β_{DTWR} coincided with increases in idiosyncratic volatility. β_{DTWR} still has applications from a risk management perspective, as it provides an intuitive and novel way of adjusting for sluggishness in the incorporation of market-wide information into stock prices. Importantly, this result corroborates Liu et al. (2018) who address the relation between beta and idiosyncratic volatility. Given the relation between β_{DTWR} and idiosyncratic volatility, we use bootstrapping to correct for bias in the estimation of DTW beta, as in Eq. (14).

<Insert Table 4 about here>

We calculate the average Spearman rank correlation between stock level estimates of β and $\beta_{Sw}, \beta_{DIM}, \beta_{DTW}$ in Table 4. Using New York Stock Exchange (NYSE) market capitalization breakpoints we divide our sample into micro (bottom 20%), small (middle 30%) and large (top 50%). We then calculate the correlation of $\beta_{Sw}, \beta_{DIM}, \beta_{DTW}$ with β across different time windows. Across all time horizons and size splits, we find that the correlation with β is relatively stable for the other beta measures. The primary differences we observe are that the average correlations for each measure are lower for micro stocks, and that β_{DTW} in general has a lower correlation with β than β_{DIM} and β_{SW}. This result is most pronounced for micro stocks, which suggests that β_{DTW} is most different to β for micro stocks. This supports the idea that β_{DTW} has the most influence for micro stocks which lag behind the market and we expect β for these stocks to be more downwards biased.

<Insert Table 5 about here>

Table 5 shows the time-series average of the median characteristic value within each month for decile portfolios produced by sorting from low to high β_{DTW}. Across several characteristics we find a "U-shaped" pattern, where portfolio one and portfolio ten have high values, and portfolios two to nine have relatively lower values. This is of particular prominence for ILLIQ, where we see that portfolio one and portfolio ten are both comprised of illiquid stocks. This suggests that some small, illiquid stocks have their true beta underestimated using the standard beta measure, and DTW beta finds that the betas should be substantially higher. This result supports our initial hypothesis that some small, illiquid stocks may have a downwards biased beta due to sluggishness in incorporating market-wide information into their price. This result also suggests that β_{DTW} doesn't just estimate higher β_{DTW} for all small, illiquid stocks which have low values of β but only for those stocks where the beta should be genuinely higher. β_{DTW} still finds that some small, illiquid stocks are genuine low beta stocks.

<Insert Table 6 about here>

Table 6 presents the equal-weighted (EW) and value-weighted (VW) excess returns associated with decile portfolios using univariate sorts of $\beta, \beta_{SW}, \beta_{DIM}$ and β_{DTW}. There is a statistically insignificant positive excess return associated with the High minus Low (HL) portfolio for β_{DTW}, whilst the three other metrics have statistically insignificant negative excess returns. Whilst β_{DTW} recovers portfolio returns which are increasing from low to high, the pattern is not monotonic. We find that the portfolio nine has a higher average excess return than the High portfolio for β_{DTW}. This suggests that whilst β_{DTW} is a significant improvement over the other beta measures at recovering the CAPM predicted risk-return relationship, the result is still not fully in line with the predictions of the CAPM.

<Insert Table 7 about here>

We extend our study of the β_{DTW} factors from the univariate portfolio setting, to the multivariate setting using Fama and MacBeth (1973) regressions. Table 7 presents the average coefficients from running monthly Fama and MacBeth (1973) regressions. We run four regressions on the measures of beta: β_{DTW}, β, β_{DIM}, and β_{SW}. β , in regression (ii) exhibits a statistically insignificant negative coefficient, the manifestation of the beta anomaly. β_{DIM}, and β_{SW} in regressions (iii) and (iv) exhibit statistically insignificant coefficients. β_{DTW} in regression (i) exhibits a statistically insignificant positive coefficient of 0.19, reversing the negative coefficient we found in the β case. We introduce a series of control variables to our regression, and run four regressions on each of the beta factors. Our first observation is that by introducing our set of control factors, the coefficients on β, β_{DIM}, and β_{SW} are now all positive but statistically insignificant. This result is in line with Liu et al. (2018), where they document that the beta anomaly arises from idiosyncratic volatility. Our second observation is that the specification of β_{DTW} finds a statistically significant, at the 1% confidence level, average slope coefficient when controlling for several other well-known asset pricing factors. This suggests that the excess returns associated with the β_{DTW} factor are not explained by the set of control factors we have included, and in fact the interaction with these control factors strengthens β_{DTW}.

<Insert Table 8 about here>

Having established that we can recover positive excess high-minus-low (HL) returns with our β_{DTW} measure, we examine the factor loadings and risk-adjusted alphas of HL portfolios for the four measures of beta in Table 8. We regress HL portfolio returns for β_{DTW}, β, β_{DIM}, and β_{SW} against five different risk models. We select the Fama-French three-factor model (Fama and French, 1993), Fama-French-Carhart four-factor model (Carhart, 1997), Fama-French five-factor model (Fama and French, 2015), Fama-French six-factor model (Fama and French, 2018), and the FFC4 model including FMAX (Bali et al., 2011). What we find as we move between the five risk models is that the alphas become less negative and the statistical significance weakens. For β_{DTW} under the FFC4+FMAX models we find statistically significant and positive alphas, whilst for the other three betas the presence of the FMAX factor renders their alphas statistically insignificant.

<Insert Table 9 about here>

To further examine the relationship of the beta measures, we perform dependent portfolio decile sorts between a set of control characteristics (ILLIQ, SIZE, IVOL, MOM, MAX, and BM) and the four beta measures. Table 9 presents the results of these dependent portfolio sorts. Again, we find the result that the FFC4+FMAX alphas for β_{DIM} under different control variables is statistically significant and positive for all control's except MOM.

<Insert Table 10 about here>

Table 10 presents portfolio returns, FFC4 alphas and FFC4+FMAX alphas across different size buckets and time-windows. Within each size bucket, we calculate HL portfolio returns for the four measures of beta. We then calculate the FFC4 and FFC4+FMAX alphas of these portfolio returns, and average these across different time windows. The primary finding is that in the January 1, 2000 to December 31, 2019 period the positive relationship that we recover for β_{DTW} inverts. In fact, for all measures of beta across all size buckets, the portfolio returns and alphas are more negative in the 2000/01 to 2019/12 period when compared to the 1927/07 to 1999/12 period. One potential explanation here is the reflexive nature of markets. As market participants learnt about the beta anomaly, there could have been more market flows into low beta stocks which subsequently results in their outperformance. In this case, correcting for asynchronicity is unable to resolve the beta anomaly as it is not being induced by asynchronicity between stock returns and market returns.

Our results show that by accounting for dynamic asynchronicity between stock returns and market returns, we can obtain better estimates of beta. We emphasize that we do not assume that the predictions of CAPM are correct, or that beta as a factor should generate consistently positive statistically significant excess returns. Rather, our approach allows for a better estimate of beta where stock returns and market returns are better aligned than the common approaches of Dimson (1979) and Scholes and Williams (1977). There are numerous other considerations for CAPM beta's in general, such as the fact that the beta anomaly itself has likely been arbitraged away in the past 20 years with the advent of smart beta products and investment managers actively targeting the beta anomaly. This clouds the picture for expected results for beta and whether it is indeed a risk premia. However, the use of DTW to account for asynchronicity allows for the disentanglement of effects associated with beta. For example, there have been several recent studies which examine the impact of beta and risk, such as, when a stock is added to an index it's beta with respect to that index increases (Barberis et al., 2005), ETF's increasing co-movement and betas (Da and Shive, 2018), and that increased algorithmic trading increases market co-movement and betas (Malceniece et al., 2019; Park and Wang, 2020). DTW potentially allows one to disentangle whether increases in co-movement are due to a true change in systematic risk, or are driven by a better alignment of stock returns between the benchmark and constituents and thus the stock's beta is now simply being more accurately measured.

4. Global markets price discovery

If global markets were perfectly integrated we would expect that securities which share commonality to exhibit no lead/lag in response to information flows. Lo and MacKinlay (1990) highlight that market frictions slow the transmission of information across markets, and that these market frictions can induce dynamic asynchronicity in markets. An example of this is in the case of cross-listed securities. If the exchanges on which the security is listed are all concurrently open and a piece of relevant information is released, there will be impounding of this information into the price occurring across all exchanges. The speed at which this occurs is the driver of common price discovery models. Thus, price discovery models are explicitly dealing with asynchronicity in how information is impounded into markets. However, common price discovery models typically don't account for changes in the asynchronicity within the estimation window, which can lead to errors in inference when using these models. We use DTW to measure the intraday lead/lag, at each time-step within the estimation window, between assets which share some commonality. In this study, we focus on U.K. and U.S. equity futures.

The link between the U.S. and U.K. equity markets is well studied, with several studies identifying bi-directional information transmission. Eun and Shim (1989) explore the transmission of information across global stock markets using a vector autoregression framework. They demonstrate the dominant influence of the U.S. stock market on other global markets, as well as bi-directional information transmission between the U.K. and the U.S. There is an extensive literature which explores volatility spillover effects that occur between different markets. Antonakakis et al. (2016) explore dynamic volatility spillovers between the U.K. and U.S. futures markets, and identify that the relationship is bi-directional, and that volatility in U.K. futures are net receivers of shocks to futures volumes. The literature on the U.K. and U.S. futures lead/lag relationship is relatively sparse, with the literature instead focused on the lead/lag relationship between index futures and the underlying equity market indices in the U.K. and U.S., separately. We contribute to the literature by applying DTW to measure the intraday lead/lag that exists between U.K. and U.S. futures, and demonstrate the time-varying nature of the lead/lag.

4.1. Data and method

We use intraday quotes on the CME E-mini S&P 500 futures (E-mini) and ICE FTSE 100 (FTSE 100) index futures to study the bi-directional transmission of information, to identify, and quantify any lead/lag structures which exist. It is well established that the E-mini futures are one of the most frequently traded and most liquid instruments globally, and thus one would expect this instrument to be the dominant instrument from a price discovery perspective. The E-mini futures and FTSE 100 futures, whilst not perfectly correlated, share a high amount of similarity as they reflect underlying large developed equity markets which both respond to globally relevant macroeconomic information flows. It is this link which we seek to study, and we apply DTW to draw out new insights.

We use the continuous futures chains from Refinitiv, for the CME E-mini S&P 500 futures (RIC: ESc1) and the ICE FTSE 100 index futures (RIC: FFIc1). For each day between November 13, 2001 and June 30, 2020, we take the one-second bid and ask quotes for each instrument and calculate the mid-quote. We also obtain the traded volume of each contract within the one-second interval. We use this to calculate the E-mini volume share, as the ratio of the E-mini traded volume to the combined E-mini and FTSE 100 traded volume. We remove any observations where the bid is greater than the ask, and take the previous valid quote. We calculate the log-return in each one-second interval using the mid-quote, and remove any observations in which the absolute one-second log-return is greater than 25%. For each day we determine the start and end time for each time-series as the time at which the first and last trade in either the E-mini contract or the FTSE 100 contract occurs. We apply DTW to each daily time-series of E-mini and FTSE 100 one-second log-returns, using a DTW window of 60 seconds across all time periods. To reduce the effect of the bias in the DTW algorithm, we run two iterations of the DTW algorithm. In the first run, the first time-series used in the DTW algorithm is the E-mini contract and the second time-series is the FTSE 100 contract, and in the second run the first time-series is the FTSE 100 contract and the second time-series is the E-mini contract. We obtain for each day in our sample, two estimates in each one-second interval of the lead/lag between the E-mini and FTSE 100 contracts, we then take the simple average of these two estimates, such that a positive value indicates the E-mini contract leads and a negative value indicates the FTSE 100 contract leads.

We use UTC time to synchronize data from Refinitiv, we acknowledge that there is likely a time lag in Refinitiv receiving information from the various exchanges, however this is deemed to have minimal impact at the one-second time-scale we run our analysis at. We note that there are several changes in the trading hours of the ICE FTSE 100 futures within our sample period. Prior to June 2, 2008, the FTSE 100 contract traded between 8am–5:30pm London time. Between June 2, 2008 and October 4, 2010, the FTSE 100 contract traded between 8am–9pm London time. Post October 4, 2010, the FTSE 100 contract traded between 1am–9pm London time. We also note that on November 17, 2014, ICE transitioned the FTSE 100 futures contracts from LIFFE, at this point the FTSE 100 contract traded between 7am–9pm London time. On October 1, 2015, ICE transitioned the FTSE 100 futures trading hours to 1am–9pm London time.

4.2. Results and discussion

<Insert Figure 5 about here>

We first explore the results of applying the standard Hasbrouck (1995) Information Share (IS) model to measure price discovery. Figure 5 presents the average of the upper and lower bounds estimated using the Hasbrouck IS method, where we use the one-second log-returns of the E-mini and FTSE 100 contracts. We observe that prior to 2010, the E-mini contract was where the majority of price discovery action occurred, however post-2010 we see that the share of price discovery converges to approximately equal between both instruments. This coincides with an increase in the correlation between the two instruments, which is potentially driven by the extension of the trading hours in the FTSE 100 contract that occurred in 2009 and 2010, and is also likely a function of a general increase in global co-movement in assets (Rua and Nunes, 2009). Our key reason for presenting this result, is to demonstrate the standard price discovery models generally only provide summary estimates of the lead/lag, often at the daily level. Whilst it is possible to run the models over subsets of the day, there is a limit to how fine the estimation windows can be made before model error becomes too large. One advantage of using DTW is that it is agnostic to the estimation window, and produces an estimate of the lead/lag for every observation within the estimation window.

<Insert Figure 6 about here>

Figure 6 presents the yearly average of the daily median lead/lag between the E-mini and FTSE 100 contracts, and the E-mini volume share of the E-mini and FTSE 100 futures trading volumes. We divide the trading day into six distinct periods, based on the operating phase of the two underlying equity markets. We abbreviate each of the six periods as Xy where X can take the values L or N denoting LSE and NYSE, respectively, and y with values p, o, or c denoting pre-open, open and closed, respectively. There are several intuitive results from this figure. The first of these is the compression of the lead/lag values towards zero. In particular, we see in the $LoNo$ period that as the E-mini volume share started to increase from 2006, the lead that the FTSE 100 contract had over the E-mini contract reduces close to zero. This is an expected result, as we would expect that with a higher level of trading volume, that the prices of the two contracts are going to react faster to new information, and thus the lead/lag will reduce at a higher rate. Another interesting result is the trend towards zero in the $LoNo$ period, despite the E-mini volume share remaining relatively constant across this period. This result is potentially a function of overall increases in the volume of contracts traded, as well as the overall increase in the speed of information transmission between London and Chicago that has occurred as the infrastructure used for transmitting information has improved. Finally, we observe in the $LcNc$ period, the E-mini lead remains above zero, indicating that the E-mini contract tends to be leading in this period. This result is expected, as the E-mini contract is the favored contract when the underlying equity markets are closed, and is the predominant contract used to express underlying views on overnight macroeconomic news whilst equity markets are closed, due to it's relatively higher levels of liquidity. We also note that whilst the E-mini volume share tends to be close to 100% in this period, the overall volume of contracts traded is relatively thin compared to the rest of the day.

<Insert Table 11 about here>

Table 11 presents summary statistics on our lead/lag estimates, divided into the periods of: 2001–2010, 2011–2020, and 2001–2020. Note that there are no entries in Panel A for $LcNc$, as before 2009 FTSE 100 futures contract did not commence trading until the LSE had opened at 8am London time. First, we confirm that the periods of $LcNc$ exhibit a higher standard deviation, particularly when compared to the $LoNo$ period. Our estimates of DTW show that when comparing Panel A and Panel B, that there has been an overall decrease in the standard deviation of the lead/lag, as well as an overall decrease in the level of the lead/lag (as shown by the mean). Overall, these summary statistics present reasonable results which support expectations about the relationship between the E-mini and FTSE 100 contracts.

One challenging result which Table 11 presents is the $LcNo$ period, in which the average lead/lag suggests that the FTSE 100 contract is leading the E-mini contract, whilst the median lead/lag suggests the inverse. What we observed is that during this period the lead/lag tends to be approximately zero except around the period when the LSE is closing. When the LSE closes, there is a strong impulse response of the lead/lag in which the FTSE 100 contract takes over, followed by a reversion to zero in the lead/lag. This explains why we see a mean which suggests that FTSE 100 leads, whilst the median suggests the E-mini leads. Having examined the time-series behavior of the average daily behavior, we now focus on the intraday dynamics of the lead/lag.

Figure 7 presents the cross-sectional average of the intraday lead/lag between the E-mini and FTSE 100 contracts, measured at a one second frequency, in two distinct periods: 2014–2019 and 2020, using the months January through May in each year. We use the months January through May to focus on the 2020 COVID-19 market events. We show the shift in lead/lag behavior which has occurred in 2020 as a result of the COVID-19 pandemic, and the elevated levels of both market volatility and market activity which have been seen. The 2014–2019 lead/lag (solid black line) shows several interesting features. The most common thematic is that lead/lag undergoes impulse-like responses to significant intraday events such as market open/closes and shifts in relative levels of liquidity that occur at periods such as the US economic announcements that occur at 1:30pm UTC. This behavior is likely a result of the sudden sharp shifts that occur in the trading of the futures contracts as the volumes in the underlying equity markets shift, and the relevance of the information being impounded into underlying equity prices shifts between markets.

<Insert Figure 7 about here>

Table 12 presents a narrative to highlight some of the intraday insights that our DTW results allow. We focus on the 2014-2019 result.

<Insert Table 12 about here>

Table 12 focuses on the 2014–2019 results (using January through May months). We contrast this result against the 2020 lead/lag (solid red line) in Figure 7. Our result shows that the E-mini contract tends to be a lot more dominant in the periods in which the NYSE is closed, compared to the 2014–2019 period. In particular we see that the E-mini volume share is elevated in the period between 7am–2:30pm UTC. An interesting observation is that over the 8am–2:30pm period, we see that the difference in the E-mini volume share in the 2020 period and 2014–2019 period actually compresses, and this is reflected in the lead/lag estimates as well. We see that the overall differential between the 2014–2019 lead/lag and 2020 lead/lag gets smaller as we move from 8am–2:30pm. We then see that after 2:30pm the behavior of the lead/lag is quite similar in the two samples.

<Insert Figure 8 about here>

Figure 8 shows the 2002–2009 period contrasted against the 2014–2019 period. In a similar vein to Figure 7 we see distinct periods of lead/lag behavior. We note that in the 2002–2009 period, the futures were typically only able to be traded between 6am UTC and 5:30pm UTC. We find that between the two time periods, the overall lead/lag behaviors are similar, however the magnitude of the lead/lag is what varies. In the 8am–2:30pm UTC phase, we see in the 2002–2009 period that the lead/lag starts to move closer to zero after 12pm UTC (when the NYSE pre-open phase begins) and we also see some swings in the lead/lag around the 1:30pm UTC US macroeconomic news announcements. Overall, we see that the lead/lag in the 2002–2009 period tend to be larger than that of the 2014–2019 period. This result is expected, however the intraday variation in these differentials is not constant and is worth further investigation. In particular, we see that in the 8am–2:30pm period, the difference in the lead/lag between 2002–2009 and 2014–2019 is substantially larger than in the 2:30pm–4:30pm UTC period. This is likely a function of both equities markets being fully open for trading, and thus information is being processed a lot quicker in both contracts.

5. Conclusion

The presence of dynamic asynchronicity is an important consideration when modeling financial time-series. Asynchronicity can occur at all frequencies of time-series in financial settings. From the quarterly level lead/lag of macroeconomic information and asset returns, to the millisecond lead/lag between equity futures traded in Chicago and ETF's traded in New York. We present DTW as a method to not only account for asynchronicity between financial time-series, but also to account for the time-varying nature of asynchronicity.

Using bootstrapped simulations, we validate the use of DTW for measuring asynchronicity in financial time-series. DTW proves to be robust at capturing different forms of induced lead/lag between financial time-series, and we uncover insights into the influence of volatility and noise on DTW's ability to recover lead/lag structures. To further demosntrate DTW's usefulness, we explore two empirical settings at different time frequencies. We use DTW to provide a better measure of beta which accounts for dynamic asynchronicity between stock returns and market returns, helping in resolving the beta anomaly. Our approach successfully recovers a positive relationship between high beta stocks and excess returns. Using DTW we are able to measure the intraday dynamics in lead/lag structures between E-mini and FTSE 100 futures, and demonstrate the time-varying behavior in the price leadership.

Through our empirical studies, we demonstrate situations in which DTW is able to generate new and interesting insights into traditional financial economics problems. DTW has scope for application to other important problems where dynamic asynchronicity in financial time-series can drive errors in inference in traditional empirical models.

Appendix A: Variable definitions

Beta estimation

- Beta (β): We follow Fama and MacBeth (1973) and measure beta by regressing excess stock returns on a constant and the market excess return:

$$r_{i,t} - r_{f,t} = \alpha_i + \beta_i(r_{m,t} - r_{f,t}) + \epsilon_{i,t}, \tag{8}$$

  $r_{i,t}$ is the return of stock i during period t, $r_{f,t}$ is the risk-free rate during period t, $r_{m,t}$ is the return of the market portfolio during period t, and $\epsilon_{i,t}$ is the regression residual of stock i during period t. We measure each beta during period t, using the daily return data from the 12-month period from months $t-11$ through t, inclusive. We require a minimum of 200 observations for the regression to be valid, and this applies to all measures of beta using daily data, described below.

- Dimson beta (β_{DIM}): We follow Dimson (1979) and estimate a beta that is designed to adjust for potential infrequent trading events. We add five days of lagged market returns and five days of forward market returns into the regression for beta as follows:

$$r_{i,t} - r_{f,t} = \alpha_i + \beta_i^{(k)} \left(\sum_{k=-5}^{5} r_{m,t-k} - r_{f,t-k} \right) + \epsilon_{i,t} \tag{9}$$

  The Dimson beta estimator is then the sum of the estimated beta coefficients from Eq. (9):

$$\beta_{DIM_i} = \sum_{k=-5}^{5} \beta_i^{(k)} \tag{10}$$

- Scholes-Williams beta (β_{SW}): We follow Scholes and Williams (1977) to calculate an adjusted beta. We use three regressions to obtain three measures of betas. The first regression uses contemporaneous market returns, as in Eq. (8). The second regression uses the one day lagged market excess return as the explanatory variable (Eq. (11)), and the third regression uses the one day forward market excess return as the explanatory variable (Eq. (12)).

$$r_{i,t} - r_{f,t} = \alpha_i + \beta_i^{-}(r_{m,t} - r_{f,t}) + \epsilon_{i,t}^{-} \tag{11}$$

$$r_{i,t} - r_{f,t} = \alpha_i + \beta_i^{+}(r_{m,t} - r_{f,t}) + \epsilon_{i,t}^{+} \tag{12}$$

The Scholes-Williams beta estimator is given by:

$$\beta_{SW_i} = \frac{\beta_i^{-} + \beta_i + \beta_i^{+}}{1 + 2\rho} \tag{13}$$

Where ρ is the first order autocorrelation of the market excess return over the estimation window.

- DTW beta and DTW t-statistic (β_{DTW} and β_{DTWT}): We measure β_{DTW} for each stock i during month t as follows:

  1. Start with daily stock excess returns and daily market excess returns over the months $t - 11$ through t, inclusive
     (a) (Bootstrapping step only): Randomly permute the daily stock excess returns
  2. Calculate the compounded cumulative returns series for both the stock and the market over the estimation period
  3. Normalize each time-series by subtracting the mean and dividing by the standard deviation over the estimation period
  4. Run DTW using a window of 20 days to obtain the optimal alignment path between the normalized cumulative stock returns series and normalized cumulative market returns series
  5. Use the optimal alignment path from Step 4 to align the original daily stock excess returns and daily market excess returns from Step 1. It is worth highlighting that we do not use the compounded cumulative returns series in the estimation of beta, these are simply used to obtain the optimal alignment path and then discarded
  6. Use these optimally aligned returns series to estimate beta as in Eq. (8)

For each stock i in month t, we run a bootstrapping simulation using the above procedure. There is an additional Step 1(a) which is run for each bootstrapping simulation. We run

$k = 100$ simulations where in each simulation we measure the $\beta_{DTWR_{i,k}}$ from permuting the stock excess returns series. We calculate:

$$\beta_{DTWi} = \beta_{DTWRi} - \frac{1}{100} \sum_{k=1}^{100} \beta_{DTWR_{i,k}}, \tag{14}$$

$$\beta_{DTWTi} = \frac{\frac{1}{100} \sum_{k=1}^{100} \beta_{DTWR_{I,k}}}{\sqrt{\frac{\sum_{k=1}^{100} (\beta_{DTWR_{I,k}} - AVG(\beta_{DTWR_{i,k}}))^2}{100}}} \tag{15}$$

This bootstrapping can also be represented in standard beta notation as:

$$\beta_i = \rho_{i,m} \frac{\sigma_i}{\sigma_m}$$
$$\beta_i^* = \rho_{i,m}^* \frac{\sigma_i}{\sigma_m}$$
$$\beta_i^* = \frac{1}{M} \sum_{k=1}^{K} \rho_{i,m}^* \frac{\sigma_i}{\sigma_m}$$

β_i being the standard beta, and β_i^* being the beta obtained from permuting the stock excess returns (β_{DTWR_i}). It follows that,

$$\beta_i - \bar{\beta}_i^* = \left(\rho_{i,m} - \bar{\rho}_{i,m}^*\right) \frac{\sigma_i}{\sigma_m}$$

and theoretically, we have that $\bar{\rho}_{i,m}^*$ should be zero.

Control variables

- Book-to-market (BM): We follow Fama and French (1992) and compute a firm's book-to-market ratio in month t using the market value of its equity at the end of December in the previous year and the book value of common equity plus balance-sheet deferred taxes minus preferred stock for the firm's latest fiscal year ending in the prior calendar year.

- Idiosyncratic volatility (IVOL): We estimate idiosyncratic volatility using the Fama-French three factor model:

$$r_{i,t} - r_{f,t} = \alpha_i + \beta_{MKT,i}(r_{m,t} - r_{f,t}) + \beta_{SMB,i}(SMB_t) + \beta_{HML,i}(HML_t) + \epsilon_{i,t}, \tag{16}$$

SMB_t is the return to the book-to-market factor portfolio, HML_t is the return to the momentum factor portfolio, and $\epsilon_{i,t}$ is the idiosyncratic return of stock i on day t. The idiosyncratic volatility of stock i in month t is defined as the standard deviation of the daily residuals using return data from the 12-month period from months $t-11$ through t, inclusive.

- Illiquidity (ILLIQ): We follow Amihud (2002) and measure illiquidity for each stock in month t as the ratio of the absolute monthly stock return to its dollar trading volume:

$$ILLIQ_{i,t} = \frac{|r_{i,t}|}{VOLD_{i,t}}, \tag{17}$$

  $r_{i,t}$ is the return to stock i on day t, and $VOLD_{i,t}$ is the dollar trading volume of stock i on day t. We use data from the 12-month period from months $t-11$ through t, inclusive, to calculate $ILLIQ_{i,t}$.

- Maximum return (MAX): We follow Bali et al. (2011) and measure MAX as the maximum daily return within a month.

- Minimum return (MIN): We follow Bali et al. (2011) and measure MIN as the minimum daily return within a month.

- Momentum (MOM): We follow Jegadeesh and Titman (1993) and measure momentum for each stock in month t at the cumulative return on the stock over the previous 11 months starting two months ago.

- Short-term reversal (REV): We follow Jegadeesh (1990) and Lehmann (1990) and measure short-term reversal for each stock as the return on the stock over the previous month $t-1$.

- Size (SIZE): We follow existing literature and measure firm size as the natural logarithm of the market value of equity (price times shares outstanding in millions of dollars) at the end of month $t-1$ for each stock.

- Systematic and idiosyncratic skewness (SSKEW & ISKEW): We follow Harvey and Siddique (2000) and decompose skewness into idiosyncratic and systematic components by running the following regression for each stock:

$$r_{i,t} - r_{f,t} = \alpha_i + \beta_i(r_{m,t} - r_{f,t}) + \gamma_i(r_{m,t} - r_{f,t})^2 + \epsilon_{i,t} \tag{18}$$

$r_{i,t}$ is the return on stock i on day t, $r_{m,t}$ is the return on the market on day t, $r_{f,t}$ is the risk-free rate on day t, and $\epsilon_{i,t}$ the idiosyncratic return for stock i on day t. From this regression, ISKEW of stock i in month t is defined as the skewness of daily residuals $\epsilon_{i,t}$ from the prior 12 months. SSKEW of stock i in month t is the estimated slope coefficient $\gamma_{i,t}$.

- Total skewness (TSKEW): We compute the total skewness of stock i for month t using 12 months of daily returns within year t:

$$TSKEW_{i,t} = \frac{1}{D_t} \sum_{d=1}^{D_t} \left(\frac{r_{i,d} - \mu_i}{\sigma_i} \right) \tag{19}$$

D_t is the number of trading days in year t, $r_{i,d}$ is the return on stock i on day d, μ_i is the mean of returns of stock i in year t, and σ_i is the standard deviation of returns of stock i in year t.

References

Amihud, Y. (2002), Illiquidity and stock returns: Cross-section and time-series effects, *Journal of Financial Markets* **5**(1), 31–56.

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Vega, C. (2007), Real-time price discovery in global stock, bond and foreign exchange markets, *Journal of International Economics* **73**(2), 251–277.

Ang, A., Hodrick, R. J., Xing, Y. and Zhang, X. (2006), The cross-section of volatility and expected returns, *Journal of Finance* **61**(1), 259–299.

Antonakakis, N., Floros, C. and Kizys, R. (2016), Dynamic spillover effects in futures markets: UK and US evidence, *International Review of Financial Analysis* **48**, 406–418.

Audrino, F. and Bühlmann, P. (2004), Synchronizing multivariate financial time series, *Journal of Risk* **6**(2), 81–106.

Bali, T. G., Brown, S., Murray, S. and Tang, Y. (2017), A lottery-demand-based explanation of the beta anomaly, *Journal of Financial and Quantitative Analysis* **52**(6), 2369–2397.

Bali, T. G., Cakici, N. and Whitelaw, R. F. (2011), Maxing out: Stocks as lotteries and the cross-section of expected returns, *Journal of Financial Economics* **99**(2), 427–446.

Barberis, N., Shleifer, A. and Wurgler, J. (2005), Comovement, *Journal of Financial Economics* **75**(2), 283–317.

Black, F., Jensen, M. and Scholes, M. (1972), The capital asset pricing model: some empirical tests, *Journal of Financial Economics* **128**(1), 79–121.

Blitz, D. C. and van Vliet, P. (2007), The volatility effect, *Journal of Portfolio Management* **34**(1), 102–113.

Blitz, D., van Vliet, P. and Baltussen, G. (2019), The volatility effect revisited, *Journal of Portfolio Management* **46**(2), 45–63.

Burns, P., Engle, R. F. and Mezrich, J. J. (1998), Correlations and volatilities of asynchronous data, *Journal of Derivatives* **5**(4), 7–18.

Carhart, M. M. (1997), On persistence in mutual fund performance, *Journal of Finance* **52**(1), 57–82.

Chen, Y., Mantegna, R. N., Pantelous, A. A. and Zuev, K. M. (2018), A dynamic analysis of S&P 500, FTSE 100 and EURO STOXX 50 indices under different exchange rates, *PLOS ONE* **13**(3), 1–40.

Da, Z. and Shive, S. (2018), Exchange traded funds and asset return correlations, *European Financial Management* **24**(1), 136–168.

de Jong, F. and Nijman, T. (1997), High frequency analysis of lead-lag relationships between financial markets, *Journal of Empirical Finance* **4**(2), 259–277. High Frequency Data in Finance, Part 1.

Dimson, E. (1979), Risk measurement when shares are subject to infrequent trading, *Journal of Financial Economics* **7**(2), 197–226.

Dobrev, D. and Schaumburg, E. (2016), High-frequency cross-market trading: Model free measurement and applications, *Working Paper* .

Eun, C. S. and Shim, S. (1989), International transmission of stock market movements, *Journal of Financial and Quantitative Analysis* **24**(2), 241–256.

Fama, E. F. and French, K. R. (1992), The cross-section of expected stock returns, *Journal of Finance* **47**(2), 427–465.

Fama, E. F. and French, K. R. (1993), Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* **33**(1), 3–56.

Fama, E. F. and French, K. R. (2015), A five-factor asset pricing model, *Journal of Financial Economics* **116**(1), 1–22.

Fama, E. F. and French, K. R. (2018), Choosing factors, *Journal of Financial Economics* **128**(2), 234–252.

Fama, E. F. and MacBeth, J. D. (1973), Risk, return, and equilibrium: Empirical tests, *Journal of Political Economy* **81**(3), 607–636.

Franses, P. H. and Wiemann, T. (2020), Intertemporal similarity of economic time series: An application of dynamic time warping, *Computational Economics* **56**(1), 59–75.

Frazzini, A. and Pedersen, L. H. (2014), Betting against beta, *Journal of Financial Economics* **111**(1), 1–25.

Gonzalo, J. and Granger, C. (1995), Estimation of common long-memory components in cointegrated systems, *Journal of Business & Economic Statistics* **13**(1), 27–35.

Harvey, C. R. and Siddique, A. (2000), Conditional skewness in asset pricing tests, *Journal of Finance* **55**(3), 1263–1295.

Hasbrouck, J. (1995), One security, many markets: Determining the contributions to price discovery, *Journal of Finance* **50**(4), 1175–1199.

Hayashi, T. and Yoshida, N. (2005), On covariance estimation of non-synchronously observed diffusion processes, *Bernoulli* **11**(2), 359–379.

Ito, K. and Sakemoto, R. (2020), Direct estimation of lead–lag relationships using multinomial dynamic time warping, *Asia-Pacific Financial Markets* **27**(3), 325–342.

Jegadeesh, N. (1990), Evidence of predictable behavior of security returns, *Journal of Finance* **45**(3), 881–898.

Jegadeesh, N. and Titman, S. (1993), Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* **48**(1), 65–91.

Kawaller, I. G., Koch, P. D. and Koch, T. W. (1987), The temporal price relationship between S&P 500 futures and the S&P 500 index, *Journal of Finance* **42**(5), 1309–1329.

Keogh, E. and Pazzani, M. (2002), Derivative dynamic time warping, *First SIAM International Conference on Data Mining* **1**.

Lakonishok, J. and Shapiro, A. C. (1986), Systematic risk, total risk and size as determinants of stock market returns, *Journal of Banking & Finance* **10**(1), 115–132.

Lehmann, B. N. (1990), Residual risk revisited, *Journal of Econometrics* **45**(1-2), 71–97.

Li, M., Yin, X. and Zhao, J. (2020), Does program trading contribute to excess comovement of stock returns?, *Journal of Empirical Finance* **59**, 257–277.

Lintner, J. (1965), Security prices, risk, and maximal gains from diversification, *Journal of Finance* **20**(4), 587–615.

Liu, J., Stambaugh, R. F. and Yuan, Y. (2018), Absolving beta of volatility's effects, *Journal of Financial Economics* **128**(1), 1–15.

Lo, A. W. and MacKinlay, A. C. (1990), An econometric analysis of nonsynchronous trading, *Journal of Econometrics* **45**(1), 181–211.

Malceniece, L., Malcenieks, K. and Putniņš, T. J. (2019), High frequency trading and co-movement in financial markets, *Journal of Financial Economics* **134**(2), 381–399.

Martens, M. and Poon, S.-H. (2001), Returns synchronization and daily correlation dynamics between international stock markets, *Journal of Banking & Finance* **25**(10), 1805–1827.

Newey, W. K. and West, K. D. (1987), A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* **55**(3), 703–708.

Novy-Marx, R. and Velikov, M. (2022), Betting against betting against beta, *Journal of Financial Economics* **143**(1), 80–106.

Ozturk, S. R., van der Wel, M. and van Dijk, D. (2017), Intraday price discovery in fragmented markets, *Journal of Financial Markets* **32**, 28–48.

Park, A. and Wang, J. (2020), Did trading bots resurrect the CAPM?, *Working paper* .

Putniņš, T. J. (2013), What do price discovery metrics really measure?, *Journal of Empirical Finance* **23**, 68–83.

Reinganum, M. R. (1981), Misspecification of capital asset pricing: Empirical anomalies based on earnings' yields and market values, *Journal of Financial Economics* **9**(1), 19–46.

Rua, A. and Nunes, L. C. (2009), International comovement of stock market returns: A wavelet analysis, *Journal of Empirical Finance* **16**(4), 632–639.

Sakoe, H. and Chiba, S. (1978), Dynamic programming algorithm optimization for spoken word recognition, *IEEE Transactions on Acoustics, Speech, and Signal Processing* **26**(1), 43–49.

Salvador, S. and Chan, P. (2007), Toward accurate dynamic time warping in linear time and space, *Intell. Data Anal.* **11**(5), 561–580.

Scherer, B. (2013), Synchronize your data or get out of step with your risks, *Journal of Derivatives* **20**, 75–84.

Scholes, M. S. and Williams, J. (1977), Estimating betas from nonsynchronous data, *Journal of Financial Economics* **5**(3), 309–327.

Shkilko, A. and Sokolov, K. (2020), Every cloud has a silver lining: Fast trading, microwave connectivity, and trading costs, *Journal of Finance* **75**(6), 2899–2927.

Shumway, T. (1997), The delisting bias in CRSP data, *Journal of Finance* **52**(1), 327–340.

Shumway, T. and Warther, V. A. (1999), The delisting bias in CRSP's NASDAQ data and its implications for the size effect, *Journal of Finance* **54**(6), 2361–2379.

Wallace, D., Kalev, P. S. and Lian, G. (2019), The evolution of price discovery in US equity and derivatives markets, *Journal of Futures Markets* **39**(9), 1122–1136.

Yan, B. and Zivot, E. (2010), A structural analysis of price discovery measures, *Journal of Financial Markets* **13**(1), 1–19.

Zhang, Z., Tang, P. and Duan, R. (2015), Dynamic time warping under pointwise shape context, *Information Sciences* **315**, 88–101.

Table 1: Baseline DTW estimation results for four lead/lag scenarios
This table presents the results for measuring lead/lag using DTW in four simulated scenarios. We simulate two assets which share a common fundamental value, and both track the common fundamental value with some time-delay and noise. We use the parameters $N = 10,000$, $u_t = 1$, $s_{1,t} = s_{2t} = 0.5$ and $W = 60$. For each simulation, we aim to produce a true lag of ten (i..e., where p_2 lags p_1 by ten units on average across the simulation). In the constant lead/lag scenario we set $\delta_{1,t} = 5$ and $\delta_{2t} = 15$. In the gradual lead/lag scenario we set $\delta_{1,t} = 5$ and the initial $\delta_{2t} = 30$. In the oscillating lead/lag scenario we use the functional form of $\delta_{2,t} = 10\sin(2\pi x_t)$ where $x_t = \{1, 2, ..., 1000\}$, and $\delta_{1,t} = 5$. In the switching lead/lag scenario we start with $\delta_{1,t} = 5$ and $\delta_{2t} = 15$, and then randomly switch δ_{2t} to values between 0 and 20. Pearson's correlation is undefined for the constant simulation as there is no change in the true lag across the simulation.

	Lead/Lag Scenario			
	Constant	Gradual	Oscillating	Switching
Average of true lag	-10.00	-10.04	-10.00	-9.16
Average of DTW lag	-10.13	-9.97	-10.07	-9.21
DTW distance	0.84	0.85	0.84	0.86
$\widetilde{MAE}$ (absolute)	0.58	0.58	0.59	0.60
$\widetilde{MAE}$ (proportional to true lag (%))	5.86	5.83	5.86	6.50
STD. of true lag - DTW lag	0.94	0.89	0.95	0.95
Pearson correlation between true lag and DTW lag	-	0.98	0.99	0.80

Table 2: DTW estimated lead/lag between CRSP stocks and the market
Using DTW we estimate the lead/lag between daily U.S. stock returns and market returns. We use all U.S. stocks traded on all exchanges that are classified as ordinary common shares, starting in January 1950 and ending in December 2019. We split our sample into SML/MID/LGE stocks based on 30%/70% market capitalization cutoffs of stocks traded on the NYSE. We divide our sample into three time-buckets, 1950–1970, 1970–1990, and 1990–2019. Within each time-bucket and each size-bucket, this table represents the average and standard deviation of the daily cross-sectional median lead/lag.

	DTW Lead/Lag Estimate		
	SML	MID	LGE
1950–1970	0.81 ± 2.07	-0.04 ± 1.35	-0.75 ± 0.95
1970–1990	0.58 ± 1.84	0.09 ± 1.31	-0.60 ± 1.09
1990–2019	0.07 ± 1.51	-0.17 ± 1.37	-0.51 ± 0.99

Table 3: Summary statistics

This table presents summary statistics for DTW adjusted beta measures, and a standard set of control variables. For each month between July 1927 and November 2019, we calculate the cross-sectional average (Mean), standard deviation (SD), skewness (Skew), kurtosis (Kurt), number of observations (Nobs), minimum (Min), maximum (Max), median ($q_{0.5}$) and 5th ($q_{0.05}$), 25th ($q_{0.25}$), 75th($q_{0.75}$), 95th ($q_{0.95}$) percentiles. The table presents the time-series average of these cross-sectional statistic. β_{DTW} is the DTW adjusted beta, after adjusting for bias in the DTW estimator using bootstrapping. β_{DTWT} is the t-statistic associated with the bootstrapping used to calculate β_{DTW}. β_{DTWR} is the raw DTW adjusted beta. β is the beta estimated from a regression of daily excess stock returns on excess market returns using 12 months of daily returns. β_{DIM} is the beta estimated using the Dimson regression adjustment using 12 months of daily returns. β_{SW} is the beta estimated using the Scholes-Williams regression adjustment using 12 months of daily returns. MAX is the largest return over the prior month. MIN is the smallest return over the prior month. SIZE is the natural logarithm of the market capitalization. ILLIQ is Amihud's illiquidity measure. BM is the book-to-market ratio. MOM is the return from months $t-11$ to $t-1$. REV is the prior months return (short-term reversal). IVOL is the idiosyncratic volatility from a 12 month regression of daily returns. TSKEW is the total skewness, ISKEW is the idiosyncratic skewness and SSKEW is the systematic skewness calculated from daily returns over the past 12 months. Full variable definitions and formula are provided in Appendix A.

	Mean	SD	Skew	Kurt	Nobs	Min	Max	$q_{0.5}$	$q_{0.05}$	$q_{0.25}$	$q_{0.75}$	$q_{0.95}$
β_{DTW}	0.39	0.39	0.39	8.82	2276	-1.76	2.68	0.35	-0.16	0.14	0.60	1.06
β_{DTWT}	22.33	4.12	-0.02	0.42	2276	7.78	36.55	22.24	15.83	19.48	25.12	29.17
β_{DTWR}	1.08	0.57	1.13	5.48	2276	-0.51	4.85	0.99	0.34	0.68	1.40	2.13
β	0.92	0.57	0.52	0.76	2276	-0.86	3.27	0.86	0.12	0.50	1.28	1.95
β_{DIM}	1.07	0.76	0.45	3.42	2273	-2.63	5.35	1.01	0.01	0.57	1.51	2.38
β_{SW}	0.86	0.52	0.49	2.70	2276	-1.19	3.51	0.82	0.11	0.50	1.18	1.78
MAX	0.03	0.02	2.74	32.74	2024	0.00	0.23	0.03	0.01	0.02	0.04	0.06
MIN	-0.03	0.01	-1.43	5.35	2024	-0.13	0.00	-0.03	-0.05	-0.03	-0.02	-0.01
SIZE	11.37	1.61	0.45	0.04	2420	7.26	17.37	11.21	8.98	10.20	12.40	14.24
ILLIQ	1.47	3.92	10.18	243.52	1845	0.00	71.70	0.52	0.04	0.17	1.45	5.63
BM	0.78	0.61	5.24	101.83	1736	0.03	12.06	0.67	0.17	0.40	1.01	1.71
MOM	0.19	0.47	3.69	47.92	2251	-0.70	7.20	0.11	-0.31	-0.07	0.33	0.92
REV	1.70	11.03	2.28	38.36	2409	-41.53	135.30	0.70	-12.83	-4.31	6.36	19.08
IVOL	0.33	0.22	3.97	66.76	2412	0.02	3.49	0.29	0.12	0.20	0.41	0.70
TSKEW	0.46	1.10	2.22	24.02	2276	-6.55	10.97	0.34	-0.75	0.00	0.77	2.04
ISKEW	0.53	1.15	1.97	21.67	2276	-6.78	11.02	0.42	-0.81	0.05	0.87	2.22
SSKEW	-3.57	16.43	0.33	14.58	2276	-99.84	108.77	-3.15	-29.97	-12.22	5.32	21.35

Table 4: Historical correlations of asynchronicity adjusted betas with β

This tables presents the average Spearman rank correlation of $\beta_{DTW}, \beta_{DIM}, \beta_{SW}$ with β across different sample windows. β is the beta estimated from a regression of daily excess stock returns on excess market returns using 12 months of daily returns. β_{DIM} is the beta estimated using the Dimson regression adjustment using 12 months of daily returns. β_{SW} is the beta estimated using the Scholes-Williams regression adjustment using 12 months of daily returns.

	All			Large			Small			Micro		
	β_{DIM}	β_{SW}	β_{DTW}	β_{DIM}	β_{SW}	β_{DTW}	β_{DIM}	β_{SW}	β_{DTW}	β_{DIM}	β_{SW}	β_{DTW}
1927/07–2019/12	0.66	0.87	0.53	0.69	0.89	0.66	0.66	0.87	0.63	0.60	0.82	0.50
1927/07–1999/12	0.66	0.87	0.54	0.68	0.89	0.65	0.67	0.88	0.64	0.60	0.82	0.52
1927/07–1963/06	0.68	0.88	0.59	0.73	0.90	0.68	0.68	0.88	0.62	0.60	0.83	0.52
1963/07–1999/12	0.64	0.86	0.49	0.64	0.88	0.62	0.67	0.88	0.65	0.60	0.81	0.52
2000/01–2019/12	0.65	0.87	0.51	0.70	0.90	0.69	0.62	0.84	0.61	0.62	0.85	0.45

Table 5: Portfolio characteristics of stocks sorted on β_{DTW}

Each month from July 1927 to November 2019, decile portfolios are formed by sorting stocks based on DTW adjusted beta (β_{DTW}) over the previous month. This table presents the time-series average of the monthly median of characteristics within each decile portfolio.

Decile	β_{DTW}	β_{DTWR}	β	β_{DIM}	β_{SW}	BM	MAX	MIN	IVOL	ILLIQ	ISKEW	MOM	PRC	REV	SIZE	SKEW	TSKEW
Low β_{DTW}	-0.16	0.46	0.61	0.65	0.58	0.69	0.03	-0.03	0.31	0.60	0.50	0.10	13.41	0.96	10.81	-3.88	0.45
2	0.03	0.55	0.57	0.66	0.55	0.73	0.02	-0.02	0.24	0.38	0.38	0.10	18.68	0.74	11.28	-2.71	0.33
3	0.14	0.65	0.61	0.73	0.60	0.73	0.02	-0.02	0.24	0.35	0.36	0.11	20.84	0.74	11.45	-2.61	0.30
4	0.23	0.76	0.68	0.81	0.66	0.71	0.02	-0.02	0.24	0.38	0.36	0.12	21.99	0.73	11.52	-2.70	0.30
5	0.31	0.87	0.76	0.91	0.73	0.69	0.02	-0.02	0.25	0.42	0.37	0.12	22.56	0.76	11.53	-2.82	0.30
6	0.39	0.99	0.85	1.01	0.81	0.68	0.03	-0.02	0.26	0.50	0.39	0.12	22.50	0.77	11.50	-3.04	0.31
7	0.49	1.13	0.94	1.12	0.89	0.66	0.03	-0.02	0.28	0.55	0.40	0.12	21.66	0.71	11.44	-3.20	0.32
8	0.60	1.30	1.06	1.24	0.99	0.64	0.03	-0.03	0.30	0.66	0.43	0.12	20.38	0.78	11.33	-3.50	0.34
9	0.76	1.55	1.20	1.40	1.12	0.62	0.03	-0.03	0.33	0.96	0.46	0.12	18.04	0.73	11.14	-3.68	0.37
High β_{DTW}	1.06	2.05	1.45	1.70	1.34	0.58	0.04	-0.04	0.41	1.43	0.54	0.11	13.70	0.73	10.78	-4.31	0.46

Table 6: Excess returns on portfolios of stocks sorted on different beta measures Each month from July 1927 to November 2019, four sets of decile portfolios are formed by sorting stocks based on beta (β), Scholes-Williams adjusted beta (β_{SW}), Dimson adjusted beta (β_{DIM}) and DTW adjusted beta (β_{DTW}) which have been calculated using daily returns over the previous 12-months. This table reports the equal-weighted and value-weighted average monthly excess returns. We also present the average monthly excess returns associated with the high minus low portfolio. Returns are given in percentage terms, Newey-West (1987) adjusted t-statistics using 6 lags are reported in parentheses for the high minus low portfolio. ***, ** and * indicate statistical significance at 1%, 5% and 10% levels, respectively.

	Equal-Weighted				Value-Weighted			
Decile	β	β_{SW}	β_{DIM}	β_{DTW}	β	β_{SW}	β_{DIM}	β_{DTW}
Low	0.78	0.70	0.70	0.61	0.66	0.53	0.57	0.45
2	0.84	0.84	0.75	0.74	0.63	0.61	0.60	0.56
3	0.85	0.82	0.80	0.78	0.65	0.63	0.61	0.58
4	0.89	0.87	0.85	0.83	0.66	0.63	0.67	0.72
5	0.85	0.88	0.88	0.86	0.66	0.69	0.74	0.70
6	0.90	0.90	0.90	0.88	0.72	0.77	0.77	0.73
7	0.89	0.88	0.92	0.87	0.78	0.66	0.79	0.74
8	0.80	0.85	0.92	0.87	0.70	0.76	0.75	0.69
9	0.75	0.79	0.86	0.94	0.67	0.69	0.68	0.78
High	0.59	0.61	0.59	0.77	0.57	0.63	0.52	0.64
High-Low	-0.19	-0.08	-0.11	0.16	-0.09	0.10	-0.05	0.19
	(-0.89)	(-0.56)	(-0.37)	(1.05)	(-0.40)	(-0.22)	(-0.42)	(1.10)

Table 7: Firm-level cross-sectional regressions

This table presents the average coefficient estimates from monthly Fama and MacBeth (1973) cross-sectional regressions. Each month from July 1927 to December 2019, we regress excess stock returns during that month on lagged predictors including DTW beta, beta, Dimson beta, Scholes-Williams beta and several control variables, defined in Appendix A. Each row in the table reports the time-series average of the cross-sectional regression slope coefficients and their associated Newey-West (1987) t-statistics, adjusted using 6 lags, in parentheses. The R-squared value for each regression is reported in the far right column. ***, ** and * indicate statistical significance at 1%, 5% and 10% levels, respectively. Statistical significance is only indicated on beta variables.

	β_{DTW}	β	β_{DIM}	β_{SW}	MAX	MIN	IVOL	ILLIQ	ISKEW	MOM	STREV	SIZE	SSKEW	TSKEW	R^2
(i)	0.19														1.65%
	(1.27)														
(ii)		-0.09													4.05%
		(-0.81)													
(ii)			0.03												3.11%
			(0.39)												
(iv)				-0.02											3.97%
				(-0.15)											
(v)					-3.91	8.57	-1.24	0.05	0.00	1.05	-0.04	-0.14	0.00	-0.05	9.01%
					(-0.95)	(2.15)	(-3.26)	(1.54)	(0.00)	(6.73)	(-9.18)	(-5.48)	(-0.26)	(-0.71)	
(vi)	0.31				-5.37	9.44	-1.14	0.05	-0.02	1.12	-0.04	-0.14	0.00	-0.02	9.40%
	(3.01)***				(-1.80)	(3.40)	(-3.55)	(2.51)	(-0.72)	(7.21)	(-9.57)	(-5.46)	(0.09)	(-0.03)	
(vii)		0.15			-5.46	10.88	-1.09	0.07	-0.04	1.05	-0.04	-0.14	0.00	0.00	10.27%
		(1.33)			(-1.34)	(2.67)	(-3.52)	(1.82)	(-0.39)	(7.55)	(-9.78)	(-5.44)	(0.07)	(-0.31)	
(viii)			0.12		-4.74	9.46	-1.19	0.05	-0.02	1.09	-0.04	-0.13	0.00	-0.02	9.85%
			(1.61)		(-1.81)	(3.13)	(-3.50)	(2.18)	(-0.60)	(7.83)	(-9.48)	(-5.60)	(-0.07)	(-0.14)	
(ix)				0.21	-5.75	10.21	-1.09	0.06	-0.03	1.07	-0.04	-0.13	0.00	-0.01	10.14%
				(1.60)	(-1.37)	(2.45)	(-3.14)	(1.59)	(-0.39)	(8.33)	(-9.03)	(-5.57)	(-0.45)	(-0.34)	

Table 8: Factor loadings and risk-adjusted alphas for high-low portfolios

This table presents the factor loadings and risk-adjusted alphas that result from regression the LS portfolio returns based on high-low portfolio sorts of different beta measures against various asset pricing models. We regress the time-series of portfolio returns for beta (β), Dimson beta (β_{DIM}), Scholes-Williams beta (β_{SW}) and DTW beta (β_{DTW}) against the Fama-French three factor model (FF3, Fama and French (1993)), Fama-French-Carhart four factor model (FFC4, Carhart (1997)), Fama-French five factor model (FF5, Fama and French (2015)), Fama-French six-factor model (FF6, Fama and French (2018)), and FFC4 with the FMAX factor (FFC4+FMAX, Bali et al. (2011)). Each row in the table reports the regression slope coefficients and their associated Newey-West (1987) t-statistics, adjusted using 6 lags, in parentheses. The R-squared value for each regression is reported in the far right column.

Risk Model	Beta	α	Mkt-RF	SMB	HML	UMD	CMA	RMW	FMAX	N	R^2
FF3	β	-0.97	1.12	0.44	-0.13					1109	0.68
		(-6.12)	(21.65)	(2.38)	(-1.27)						
	β_{SW}	-0.83	0.94	0.49	0.03					1109	0.67
		(-5.85)	(18.56)	(4.18)	(0.28)						
	β_{DIM}	-0.94	1.11	0.57	0.02					1109	0.72
		(-6.22)	(21.18)	(4.30)	(0.19)						
	β_{DTW}	-0.39	0.61	0.37	0.23					1109	0.60
		(-3.18)	(12.24)	(3.53)	(2.19)						
FFC4	β	-0.81	1.09	0.43	-0.21	-0.17				1109	0.69
		(-5.04)	(23.11)	(2.38)	(-1.93)	(-2.33)					
	β_{SW}	-0.61	0.89	0.48	-0.08	-0.23				1109	0.69
		(-4.03)	(21.31)	(4.21)	(-0.81)	(-3.21)					
	β_{DIM}	-0.71	1.06	0.56	-0.09	-0.24				1109	0.74
		(-4.79)	(23.31)	(4.30)	(-0.85)	(-3.51)					
	β_{DTW}	-0.16	0.56	0.36	0.11	-0.24				1109	0.63
		(-1.25)	(14.24)	(3.46)	(1.28)	(-4.44)					
FF5	β	-0.45	0.93	0.42	-0.22		-0.75	-0.77		678	0.76
		(-2.37)	(13.16)	(5.18)	(-1.90)		(-4.23)	(-5.67)			
	β_{SW}	-0.36	0.73	0.44	-0.09		-0.69	-0.66		678	0.73
		(-2.52)	(11.32)	(6.38)	(-0.85)		(-4.39)	(-4.94)			
	β_{DIM}	-0.38	0.91	0.45	-0.19		-0.73	-0.79		678	0.77
		(-2.2)	(13.09)	(6.10)	(-1.66)		(-4.45)	(-6.07)			
	β_{DTW}	0.10	0.35	0.24	-0.07		-0.40	-0.59		678	0.58
		(0.86)	(7.52)	(4.09)	(-0.86)		(-3.21)	(-4.97)			
FF6	β	-0.33	0.90	0.43	-0.31	-0.17	-0.7	-0.72		678	0.76
		(-1.73)	(13.54)	(5.62)	(-2.97)	(-2.76)	(-4.35)	(-6.36)			
	β_{SW}	-0.20	0.69	0.46	-0.21	-0.23	-0.62	-0.61		678	0.75
		(-1.34)	(12.06)	(7.18)	(-2.66)	(-3.36)	(-4.99)	(-5.36)			
	β_{DIM}	-0.24	0.88	0.47	-0.29	-0.20	-0.66	-0.75		678	0.78
		(-1.36)	(13.65)	(6.74)	(-3.15)	(-3.03)	(-4.77)	(-6.86)			
	β_{DTW}	0.22	0.33	0.25	-0.15	-0.16	-0.35	-0.55		678	0.60
		(1.81)	(8.02)	(4.73)	(-2.40)	(-2.62)	(-3.28)	(-4.94)			
FFC4+FMAX	β	-0.09	0.65	0.12	-0.21	-0.21			0.81	678	0.84
		(-0.58)	(11.28)	(1.64)	(-3.24)	(-4.50)			(12.67)		
	β_{SW}	0.00	0.48	0.19	-0.14	-0.27			0.69	678	0.83
		(0.02)	(10.28)	(3.31)	(-2.34)	(-6.96)			(14.16)		
	β_{DIM}	-0.01	0.63	0.17	-0.19	-0.24			0.80	678	0.85
		(-0.06)	(11.75)	(2.61)	(-2.85)	(-6.21)			(14.84)		
	β_{DTW}	0.31	0.19	0.12	-0.08	-0.19			0.45	678	0.65
		(2.76)	(5.66)	(2.19)	(-1.57)	(-4.78)			(7.88)		

Table 9: Bivariate dependent sorts of measures of beta

This table presents the results of bivariate dependent portfolio sorts between measures of beta and stock returns after controlling for selected control characteristics. For each month, all stocks are sorted into 100 portfolios, first based on the control variable and then on the beta measure. We present the time-series means of equal-weighted excess returns of the average control variable decile portfolio within each beta decile portfolio. We also present the mean return differences of the high-minus-low portfolios (H-L) and CAPM alphas, Fama-French three-factor alphas (FF3), Fama-French-Carhart four-factor alphas (FFC4), Fama-French five-factor alphas, Fama-French five-factor and UMD alphas (FF6) and FFC4 and FMAX alphas (FFC4+FMAX) for the H-L portfolios. Newey-West (1987) t-statistics, adjusted using 6 lags, are reported in parentheses. ***, ** and * indicate statistical significance at 1%, 5% and 10% levels, respectively.

Control	Beta	Low	2	3	4	5	6	7	8	9	High	H-L	CAPM α	FF3 α	FFC4 α	FF5 α	FF6 α	FFC4+FMAX α
ILLIQ	β	0.72	0.85	0.86	0.84	0.84	0.88	0.92	0.75	0.74	0.60	-0.12	-0.82	-0.85	-0.68	-0.24	-0.16	0.13
													(-5.05)***	(-6.00)***	(-4.74)***	(-1.59)	(-0.99)	(1.14)
	β_{SW}	0.67	0.80	0.87	0.86	0.93	0.87	0.80	0.87	0.76	0.58	-0.08	-0.76	-0.79	-0.57	-0.16	-0.03	0.24
													(-4.91)***	(-5.76)***	(-4.01)***	(-1.11)	(-0.21)	(2.14)**
	β_{DIM}	0.66	0.77	0.82	0.86	0.88	0.83	0.90	0.89	0.81	0.58	-0.09	-0.66	-0.69	-0.45	-0.20	-0.05	0.19
													(-4.58)***	(-5.17)***	(-3.28)***	(-1.51)	(-0.39)	(1.67)*
	β_{DTW}	0.62	0.73	0.78	0.84	0.85	0.84	0.87	0.84	0.88	0.72	0.10	-0.31	-0.37	-0.15	0.13	0.20	0.31
													(-2.68)***	(-3.19)***	(-1.28)	(1.24)	(1.73)*	(2.67)***
SIZE	β	0.73	0.81	0.85	0.90	0.89	0.92	0.88	0.86	0.75	0.58	-0.15	-0.90	-0.93	-0.76	-0.38	-0.25	0.00
													(-5.52)***	(-6.34)***	(-5.09)***	(-2.40)**	(-1.58)	(0.03)
	β_{SW}	0.62	0.83	0.86	0.87	0.94	0.92	0.89	0.87	0.76	0.60	-0.02	-0.76	-0.81	-0.58	-0.21	-0.06	0.18
													(-4.99)***	(-5.88)***	(-4.22)***	(-1.39)	(-0.39)	(1.55)
	β_{DIM}	0.65	0.76	0.83	0.85	0.91	0.93	0.90	0.89	0.84	0.59	-0.06	-0.69	-0.73	-0.51	-0.21	-0.06	0.16
													(-4.92)***	(-5.53)***	(-3.64)***	(-1.65)*	(-0.42)	(1.41)
	β_{DTW}	0.60	0.75	0.82	0.82	0.87	0.88	0.93	0.86	0.87	0.74	0.15	-0.30	-0.36	-0.16	0.12	0.21	0.30
													(-2.55)**	(-3.11)***	(-1.34)	(1.10)	(1.81)*	(2.75)***
IVOL	β	0.88	0.82	0.84	0.80	0.85	0.83	0.82	0.77	0.78	0.70	-0.18	-0.82	-0.83	-0.71	-0.40	-0.32	-0.14
													(-6.02)***	(-6.14)***	(-5.18)***	(-2.42)**	(-1.95)*	(-0.95)
	β_{SW}	0.80	0.78	0.82	0.88	0.83	0.83	0.81	0.80	0.74	0.80	0.00	-0.65	-0.67	-0.53	-0.27	-0.18	0.00
													(-5.06)***	(-5.31)***	(-4.18)***	(-1.84)*	(-1.20)	(0.04)
	β_{DIM}	0.76	0.79	0.79	0.81	0.83	0.84	0.86	0.84	0.81	0.77	0.01	-0.55	-0.58	-0.42	-0.29	-0.17	-0.02
													(-4.64)***	(-4.96)***	(-3.28)***	(-2.40)**	(-1.38)	(-0.14)
	β_{DTW}	0.71	0.73	0.75	0.80	0.81	0.85	0.85	0.82	0.89	0.88	0.18	-0.19	-0.24	-0.08	0.07	0.15	0.22
													(-1.92)*	(-2.48)**	(-0.82)	(0.73)	(1.40)	(2.02)**
MOM	β	0.84	0.88	0.91	0.86	0.90	0.91	0.84	0.80	0.75	0.51	-0.33	-0.93	-0.92	-0.91	-0.59	-0.52	-0.37
													(-6.49)***	(-6.68)***	(-6.54)***	(-3.54)***	(-3.09)***	(-2.37)**
	β_{SW}	0.75	0.87	0.86	0.87	0.87	0.89	0.89	0.85	0.75	0.58	-0.17	-0.80	-0.82	-0.76	-0.47	-0.38	-0.25
													(-6.01)***	(-6.60)***	(-6.18)***	(-3.23)***	(-2.65)***	(-1.85)*
	β_{DIM}	0.74	0.80	0.85	0.86	0.85	0.90	0.89	0.89	0.82	0.58	-0.16	-0.69	-0.71	-0.69	-0.44	-0.36	-0.24
													(-5.51)***	(-6.12)***	(-5.14)***	(-3.61)***	(-2.89)***	(-2.10)**
	β_{DTW}	0.69	0.78	0.78	0.82	0.84	0.84	0.86	0.92	0.93	0.74	0.05	-0.28	-0.33	-0.37	-0.07	-0.07	0.00
													(-2.77)***	(-3.34)***	(-3.61)***	(-0.78)	(-0.69)	(0.02)
MAX	β	0.75	0.76	0.80	0.78	0.84	0.85	0.86	0.85	0.85	0.85	0.11	-0.39	-0.36	-0.26	-0.14	-0.05	0.07
													(-2.86)***	(-2.82)***	(-1.93)*	(-0.84)	(-0.31)	(0.48)
	β_{SW}	0.69	0.78	0.76	0.83	0.77	0.80	0.88	0.85	0.90	0.93	0.24	-0.29	-0.30	-0.14	0.01	0.11	0.22
													(-2.24)**	(-2.46)**	(-1.17)	(0.04)	(0.76)	(1.65)*
	β_{DIM}	0.72	0.67	0.83	0.75	0.84	0.82	0.85	0.87	0.90	0.92	0.20	-0.24	-0.26	-0.09	-0.12	0.02	0.10
													(-2.12)**	(-2.40)**	(-0.81)	(-0.94)	(0.13)	(0.91)
	β_{DTW}	0.65	0.69	0.71	0.78	0.78	0.86	0.86	0.90	0.95	1.00	0.36	0.05	0.00	0.16	0.20	0.28	0.30
													(0.55)	(0.02)	(1.60)	(1.99)**	(2.54)**	(2.69)***
BM	β	0.71	0.72	0.75	0.76	0.75	0.76	0.77	0.74	0.71	0.66	-0.05	-0.65	-0.62	-0.46	-0.28	-0.18	0.05
													(-3.67)***	(-3.64)***	(-2.77)***	(-1.66)*	(-1.06)	(0.35)
	β_{SW}	0.63	0.71	0.73	0.80	0.76	0.72	0.82	0.74	0.70	0.70	0.07	-0.52	-0.48	-0.30	-0.13	-0.02	0.21
													(-3.04)***	(-3.04)***	(-1.96)**	(-0.88)	(-0.11)	(1.62)
	β_{DIM}	0.58	0.66	0.73	0.77	0.86	0.79	0.80	0.81	0.78	0.55	-0.03	-0.52	-0.51	-0.31	-0.24	-0.10	0.09
													(-3.37)***	(-3.74)***	(-2.26)***	(-1.77)*	(-0.72)	(0.79)
	β_{DTW}	0.51	0.64	0.75	0.77	0.78	0.79	0.78	0.81	0.85	0.64	0.14	-0.11	-0.11	0.01	0.08	0.16	0.26
													(-0.90)	(-0.96)	(0.09)	(0.71)	(1.38)	(2.25)**

Table 10: Historical and size performance of beta measures

This table presents the results of univariate portfolio sorts on different measures of beta split into different market capitalization and time samples. We split our sample three size categories using NYSE breakpoints: bottom 20% (micro), next 30% (small) and top 50% (large). Within each size sample we conduct univariate decile portfolio sorts using different measures of beta. In Panel A we report the average return of the different between the high and low portfolios across different time windows. In Panel B we report the Fama-French-Carhart four-factor (FFC4) alphas. In Panel C we report the FFC4+MAX alphas. Newey-West (1987) t-statistics, adjusted using 6 lags, are reported in parentheses. Note we don't report the 1927/07–1963/06 in Panel C as the FMAX factor is not available.

Panel A: Portfolio Returns

	All				Large				Small				Micro			
	β	β_{DIM}	β_{SW}	β_{DTW}	β	β_{DIM}	β_{SW}	β_{DTW}	β	β_{DIM}	β_{SW}	β_{DTW}	β	β_{DIM}	β_{SW}	β_{DTW}
1927/07–2019/12	-0.19	-0.11	-0.08	0.16	-0.01	-0.02	0.06	0.09	-0.10	-0.01	-0.03	0.06	-0.37	-0.14	-0.07	0.37
	(-0.90)	(-0.56)	(-0.37)	(1.05)	(-0.03)	(-0.11)	(0.25)	(0.55)	(-0.40)	(-0.04)	(-0.12)	(0.34)	(-1.79)	(-0.78)	(-0.29)	(2.07)
1927/07–1963/06	0.03	0.16	0.20	0.33	0.04	0.16	0.22	0.23	0.05	0.24	0.02	0.15	-0.19	-0.04	0.20	0.79
	(0.10)	(0.51)	(0.52)	(1.14)	(0.11)	(0.43)	(0.53)	(0.70)	(0.12)	(0.75)	(0.05)	(0.54)	(-0.54)	(-0.13)	(0.49)	(2.09)
1927/07–1999/12	-0.18	-0.05	-0.05	0.26	0.31	0.19	0.26	0.18	-0.01	0.03	0.19	0.14	-0.42	0.00	-0.09	0.32
	(-0.60)	(-0.19)	(-0.18)	(1.74)	(0.93)	(0.67)	(0.80)	(1.04)	(-0.04)	(0.11)	(0.58)	(0.89)	(-1.40)	(-0.02)	(-0.30)	(2.03)
1963/06–1999/12	-0.18	-0.06	-0.05	0.25	0.32	0.18	0.27	0.18	-0.02	0.01	0.19	0.16	-0.44	-0.03	-0.11	0.29
	(-0.58)	(-0.24)	(-0.18)	(1.69)	(0.93	(0.65)	(0.82)	(1.06)	(-0.06)	(0.05)	(0.57)	(0.97)	(-1.47)	(-0.14)	(-0.36)	(1.87)
2000/01–2019/12	-0.61	-0.69	-0.65	-0.32	-0.69	-0.74	-0.59	-0.30	-0.50	-0.53	-0.52	-0.28	-0.63	-0.56	-0.49	-0.28
	(-1.16)	(-1.42)	(-1.20)	(-0.80)	(-1.16)	(-1.36)	(-1.01)	(-0.75)	(-0.83)	(-0.92)	(-0.86)	(-0.56)	(-1.29)	(-1.30)	(-0.99)	(-0.79)

Panel B: FFC4 Alphas

	All				Large				Small				Micro			
	β	β_{DIM}	β_{SW}	β_{DTW}	β	β_{DIM}	β_{SW}	β_{DTW}	β	β_{DIM}	β_{SW}	β_{DTW}	β	β_{DIM}	β_{SW}	β_{DTW}
1927/07–2019/12	-0.81	-0.61	-0.71	-0.16	-0.65	-0.55	-0.57	-0.30	-0.71	-0.49	-0.63	-0.20	-0.95	-0.50	-0.63	0.15
	(-5.04)	(-4.03)	(-4.79)	(-1.25)	(-3.72)	(-3.24)	(-3.39)	(-2.18)	(-3.83)	(-2.92)	(-3.69)	(-1.27)	(-4.99)	(-2.94)	(-3.93)	(0.94)
1927/07–1963/06	-0.90	-0.65	-0.73	-0.16	-0.94	-0.66	-0.74	-0.31	-0.92	-0.54	-0.97	-0.27	-1.04	-0.63	-0.57	0.48
	(-3.67)	(-2.44)	(-3.64)	(-0.89)	(-4.05)	(-2.62)	(-3.24)	(-1.69)	(-3.44)	(-2.09)	(-4.14)	(-1.23)	(-2.43)	(-1.65)	(-1.91)	(1.41)
1927/07–1999/12	-0.49	-0.27	-0.35	0.15	0.02	-0.06	-0.01	-0.06	-0.33	-0.18	-0.13	0.08	-0.76	-0.21	-0.41	0.21
	(-2.07)	(-1.60)	(-1.67)	(1.13)	(0.09)	(-0.30)	(-0.04)	(-0.38)	(-1.28	(-0.92)	(-0.56)	(0.47)	(-3.49	(-1.20)	(-1.97)	(1.40)
1963/06–1999/12	-0.49	-0.28	-0.35	0.14	0.02	-0.07	0.00	-0.05	-0.34	-0.20	-0.13	0.09	-0.78	-0.24	-0.43	0.17
	(-2.06)	(-1.68)	(-1.68)	(1.04)	(0.09)	(-0.33)	(-0.01)	(-0.36)	(-1.31)	(-1.04)	(-0.59)	(0.56)	(-3.60)	(-1.43)	(-2.09)	(1.17)
2000/01–2019/12	-1.23	-1.18	-1.25	-0.52	-1.23	-1.22	-1.15	-0.57	-1.02	-0.98	-1.00	-0.47	-1.30	-1.04	-1.15	-0.43
	(-3.64)	(-4.30)	(-4.03)	(-1.99)	(-4.03)	(-4.10)	(-3.93)	(-2.04)	(-2.64)	(-3.04)	(-2.86)	(-1.32)	(-4.11)	(-3.70)	(-4.07)	(-1.79)

Panel C: FFC4+FMAX Alphas

	All				Large				Small				Micro			
	β	β_{DIM}	β_{SW}	β_{DTW}	β	β_{DIM}	β_{SW}	β_{DTW}	β	β_{DIM}	β_{SW}	β_{DTW}	β	β_{DIM}	β_{SW}	β_{DTW}
1927/07–2019/12	-0.09	0.00	-0.01	0.31	0.54	0.39	0.54	0.31	0.28	0.28	0.40	0.31	-0.48	-0.09	-0.20	0.30
	(-0.58)	(0.02)	(-0.06)	(2.76)	(3.54)	(2.69)	(3.88)	(2.48)	(1.73)	(1.87)	(2.59)	(1.86)	(-2.99)	(-0.69)	(-1.37)	(2.42)
1927/07–1999/12	0.08	0.13	0.16	0.34	0.73	0.46	0.64	0.17	0.35	0.30	0.48	0.26	-0.32	0.10	-0.01	0.36
	(0.46)	(0.94)	(0.96)	(2.57)	(4.26)	(2.58)	(3.80)	(1.23)	(1.87)	(1.71)	(2.79)	(1.63)	(-1.74)	(0.60)	(-0.07)	(2.46)
1963/06–1999/12	0.08	0.12	0.15	0.33	0.73	0.45	0.64	0.17	0.34	0.28	0.47	0.27	-0.34	0.06	-0.04	0.32
	(0.45)	(0.83)	(0.93)	(2.46)	(4.23)	(2.52)	(3.80)	(1.24)	(1.77)	(1.55)	(2.66)	(1.71)	(-1.85)	(0.38)	(-0.20)	(2.22)
2000/01–2019/12	-0.55	-0.52	-0.52	0.08	-0.09	-0.15	0.05	0.24	0.10	-0.04	0.06	0.25	-0.91	-0.65	-0.71	0.00
	(-1.82)	(-2.27)	(-1.95)	(0.40)	(-0.37)	(-0.64)	(0.21)	(1.15)	(0.30)	(-0.16)	(0.21)	(0.80)	(-2.86)	(-2.27)	(-2.68)	(0.01)

Table 11: Summary statistics of intraday lead/lag between E-mini S&P 500 futures and FTSE 100 futures

This table presents summary statistics for the DTW estimated lead/lag between E-mini S&P 500 futures and FTSE 100 futures. For each day from November 13, 2001 to June 30, 2020, between 1:00 am UTC and 9:00 pm UTC, we calculate the daily average (Mean), standard deviation (SD), skewness (Skew), kurtosis (Kurt), number of observations (Nobs), minimum (Min), maximum (Max), median ($q_{0.5}$) and 5th ($q_{0.05}$), 25th ($q_{0.25}$), 75th ($q_{0.75}$), 95th ($q_{0.95}$) percentiles. These tables presents the time-series average of these daily statistics. Panel A presents summary statistics between November 13, 2001 and December 31, 2010. Panel B presents summary statistics between January 4, 2011 and June 30, 2020. Panel C presents summary statistics between November 13, 2001 and June 30, 2020. For each day in our sample, we use one-second mid-quote returns on the E-mini and FTSE 100 futures contracts and measure the lead/lag using DTW. We divide the day into six periods based on UTC time, and whether the equity markets underlying the E-mini futures(New York Stock Exchange) or FTSE 100 futures (London Stock Exchange) are open, closed or in a pre-open market phase. The right-most column presents the time-series average of the average E-mini volume share, which is measured as the proportion of E-mini futures volume to the combined E-mini and FTSE 100 futures volumes, calculated using the traded volume over each one-second interval. Lead/lag values are in seconds.

Panel A: 2001–2010										
	Mean	SD	Skew	Kurt	$q_{0.5}$	$q_{0.05}$	$q_{0.25}$	$q_{0.75}$	$q_{0.95}$	Volume Share
LcNc										
LpNc	-3.95	27.28	0.02	-0.84	-3.97	-5.55	-4.76	-3.12	-2.32	0.702
LoNc	-3.15	26.48	0.09	-0.82	-4.18	-44.26	-24.22	17.93	40.06	0.296
LoNp	-2.12	26.07	0.04	-0.77	-2.15	-43.21	-22.81	18.16	39.98	0.536
LoNo	0.67	20.68	-0.01	0.23	0.89	-34.15	-13.12	14.46	34.91	0.803
LcNo	-0.23	23.06	-0.05	-0.29	0.49	-38.03	-17.17	16.46	36.70	0.998

Panel B: 2011–2020										
	Mean	SD	Skew	Kurt	$q_{0.5}$	$q_{0.05}$	$q_{0.25}$	$q_{0.75}$	$q_{0.95}$	Volume Share
LcNc	2.28	23.98	-0.09	0.01	1.19	-37.05	-9.45	16.24	39.57	0.960
LpNc	-0.38	24.01	-0.01	-0.46	-0.32	-39.00	-17.66	16.81	37.89	0.727
LoNc	-1.23	22.61	0.02	-0.26	-0.55	-38.80	-17.14	14.13	36.68	0.511
LoNp	-1.46	23.00	0.02	-0.33	-0.87	-39.35	-17.90	14.36	36.90	0.666
LoNo	0.28	17.92	0.04	1.06	0.10	-30.29	-10.18	10.83	31.06	0.856
LcNo	-0.30	19.48	-0.03	0.67	0.07	-34.00	-11.68	10.82	32.97	0.994

Panel C: 2001–2020										
	Mean	SD	Skew	Kurt	$q_{0.5}$	$q_{0.05}$	$q_{0.25}$	$q_{0.75}$	$q_{0.95}$	Volume Share
LcNc	2.28	23.98	-0.09	0.01	1.19	-37.05	-9.45	16.24	39.57	0.970
LpNc	-2.16	24.13	-0.01	-0.47	-2.14	-22.34	-11.23	6.88	17.86	0.715
LoNc	-2.19	24.54	0.05	-0.54	-2.36	-41.52	-20.67	16.02	38.37	0.405
LoNp	-1.79	24.53	0.03	-0.55	-1.51	-41.28	-20.35	16.26	38.44	0.602
LoNo	0.47	19.30	0.01	0.64	0.49	-32.22	-11.65	12.65	32.98	0.830
LcNo	-0.27	21.26	-0.04	0.19	0.28	-36.00	-14.41	13.62	34.82	0.995

Table 12: Description of intraday trading sessions on the NYSE and LSE

Time (UTC)	LSE	NYSE	Label	Description
1:00am–7:00am	Closed	Closed	$LcNc$	E-mini leads, however trading volumes in both contracts are relatively thin and we see higher levels of cross-sectional volatility in the estimates.
7:00am–8:00am	Pre-open	Closed	$LpNc$	The pre-open phase of the LSE sees volumes in the FTSE 100 contracts increase relative to the E-minis, and this causes the FTSE 100 to lead. Heading into the LSE open, we see the E-mini contract take over, before the FTSE 100 takes over just before LSE open.
8:00am–12:00pm	Open	Closed	$LoNc$	FTSE 100 leads across this entire period. The lead tends to be fairly stable.
12:00pm–2:30pm	Open	Pre-open	$LoNp$	FTSE 100 continues to lead, however we see a shift occur at 12pm when the NYSE pre-open begins and at 1:30pm when US macroeconomic news is announced. We see the E-mini contract begins to take-over heading into the NYSE open.
2:30pm–4:30pm	Open	Open	$LoNo$	E-mini leads, however as both markets are open, the lead tends to be stable and close to zero. Heading into the LSE close at 4:30pm, we see the E-mini contract start to take a stronger lead, before the FTSE 100 takes over.
4:30pm–9:00pm	Closed	Open	$LcNo$	Immediately after the LSE close, we see an impulse response where the FTSE 100 contract takes a strong lead which reverts back to zero over a 15 minute period. This is likely due to impounding of information from the market close, resulting in increased trading volumes and greater asynchronicity in the two instruments. For the rest of the period, the E-mini tends to lead however the overall lead/lag is approximately 0 with variable noise around the estimate.

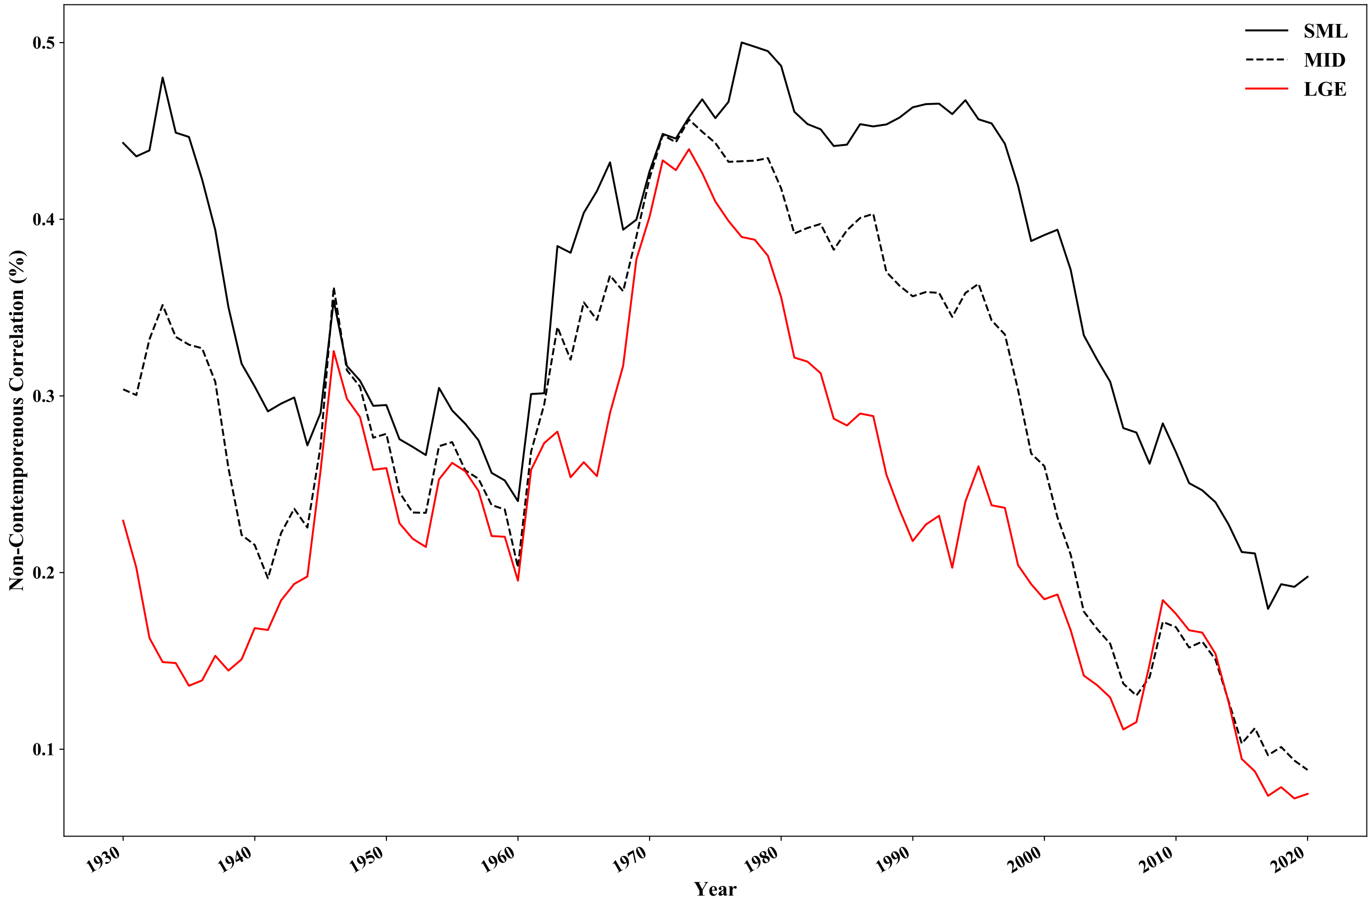

Figure 1: Proportion of non-contemporaneous correlation in CRSP sample
This figure presents the ratio of the one-day forward correlation and one-day lagged correlation between CRSP stock returns and market returns to the sum of the one-day forward correlation, contemporaneous correlation and one-day lagged correlation. Between June 1931 and December 2019, at the end of each month, using 5 years of daily data, we calculate the Pearson correlation between stock returns and market returns, as $Corr(r_{m,t}, r_{i,t+s})$ with $s = (-1, 0, +1)$. We then calculate the non-contemporaneous correlation proportion as:

$$\frac{|Corr(r_{m,t}, r_{i,t-1})| + |Corr(r_{m,t}, r_{i,t+1})|}{|Corr(r_{Mm,t}, r_{i,t-1})| + |Corr(r_{m,t}, r_{i,t})| + |Corr(r_{m,t}, r_{i,t+1})|}$$

$r_{m,t}$ is the market return for day t, and $r_{i,t+s}$is the return of stock i at day $t+s$. We separate our universe into size-buckets, SML/MID/LGE, based on 30%/70% market capitalization cutoffs of stocks traded on the NYSE. This figure presents the cross-sectional median non-contemporaneous correlation proportion at the end of each month within three size-buckets.

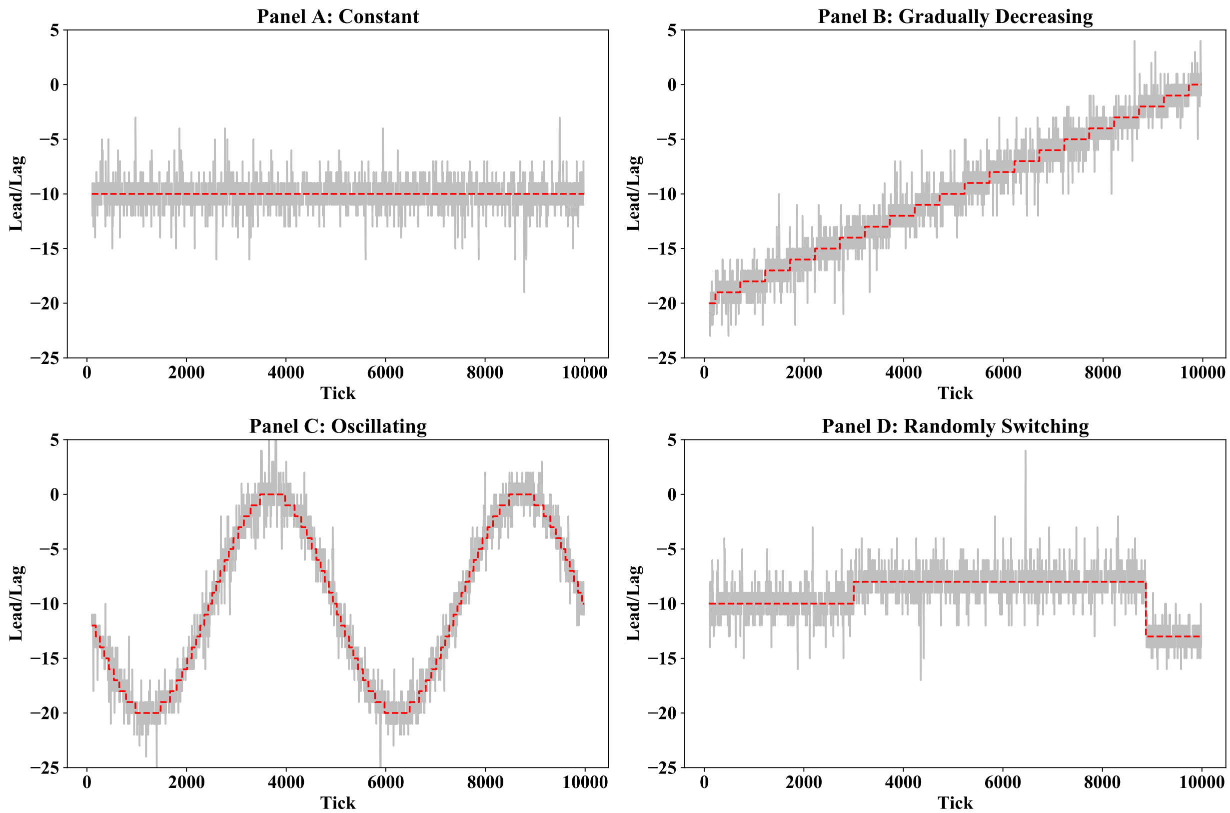

Figure 2: DTW recovery of lead/lag from four simulated scenarios
This figure presents four scenarios where an artificial lead/lag structure is induced between two simulated time-series which share a common fundamental value. DTW is used to recover this artificial lead/lag. The solid gray line is the lead/lag estimated using DTW. The dashed red line is the artificial lead/lag which is induced between the two simulated time-series.

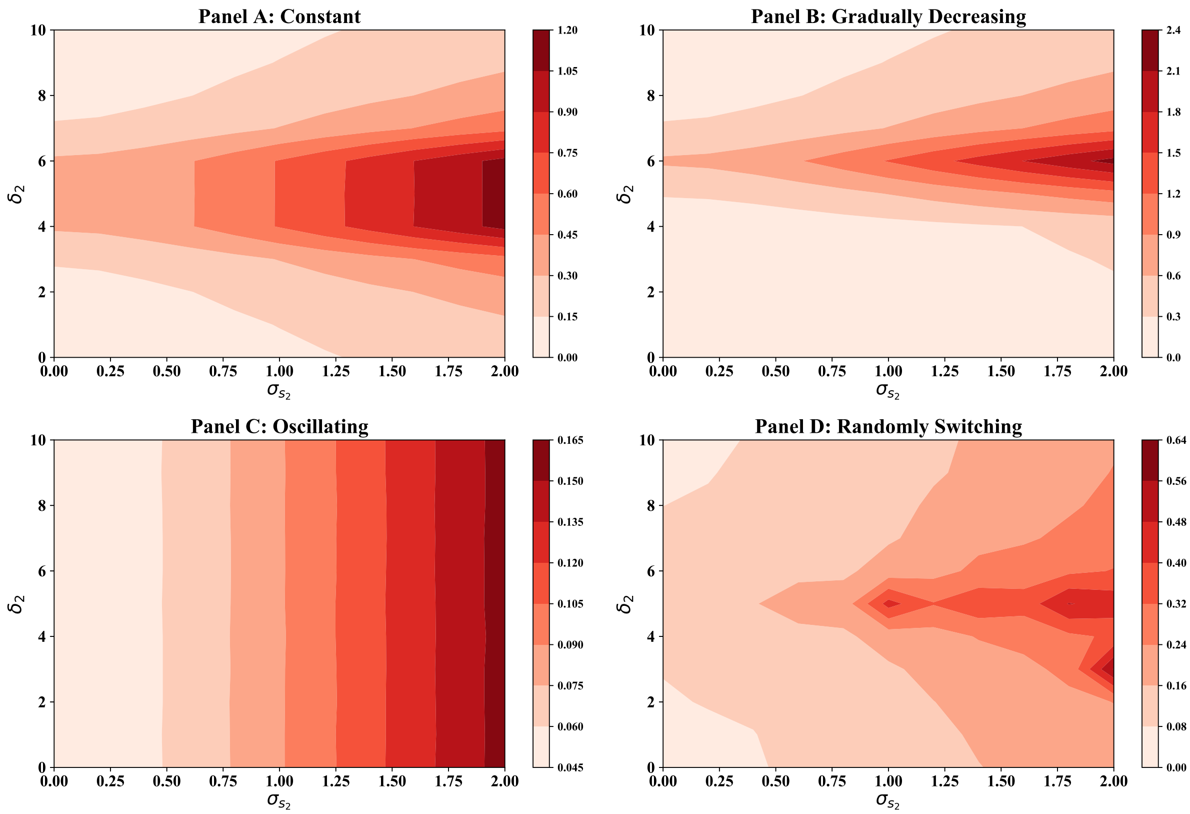

Figure 3: Effect of varying δ_2 and σ_{s_2} on accuracy of DTW estimation of lead/lag
This figure presents the mean absolute error ($\widetilde{MAE}$) between the artificially induced lag and the DTW estimated lag, when we vary σ_{s_2} and δ_2 when generating the time-series under four lead/lag scenarios. The $\widetilde{MAE}$ is calculated as the difference between the DTW estimated lag and artificially induced lag at each tick of the simulation, proportional to the average of the induced lead/lag across the simulation. We use a grid of ten values for σ_2 and δ_2. σ_2 is varied between 0 b.p. and 2 b.p. at increments of 0.2 b.p., δ_2 is varied between zero and ten at increments of one unit. For each possible pair of parameters, 1,000 simulations are run with input parameters of $N = 10,000$, $\sigma_u = 1$, $\delta_1 = 5$, $\sigma_{s_1} = 1$, $W = 60$ and $\psi = 60$. Each panel corresponds to one lead/lag scenario.

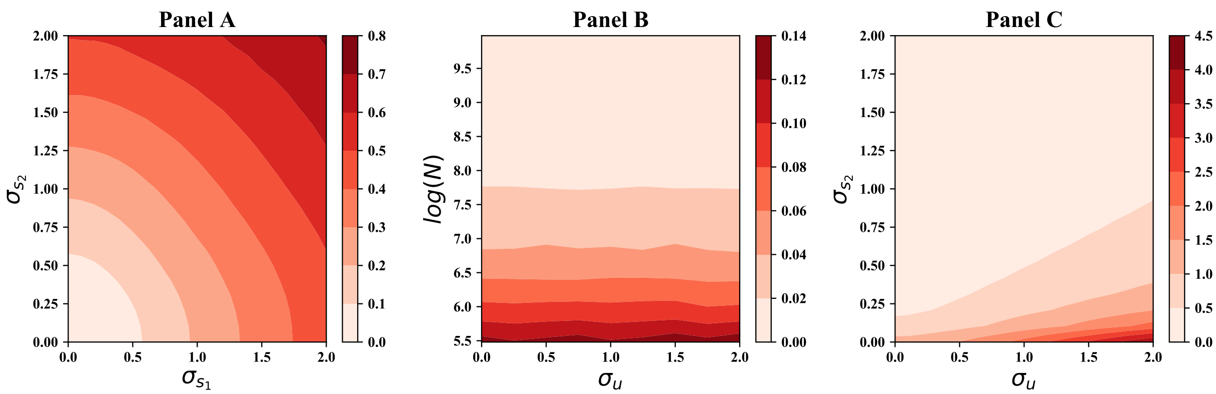

Figure 4: Effect of varying three pairs of simulation parameters on accuracy of DTW estimation of lead/lag

This figure presents the mean absolute error ($\widetilde{MAE}$) generated from varying three different parameter pairs when generating the time-series, and then estimating the lead/lag using DTW. In panel A: we vary the values of $\sigma_{s_1}$and σ_{s_2} between 0 b.p. and 2 b.p., using 0.2 b.p. increments. In panel B: we use the natural logarithm of 240, 3,600, 7,200, 14,400 and 21,600 for N, and we vary the value of σ_u between 0 b.p. and 2 b.p. using 0.2 b.p. increments. In panel C: we vary the values of σ_u and σ_{s_2} between 0 b.p. and 2 b.p., using 0.2 b.p. increments. These panels show the bootstrapped average $\widetilde{MAE}$ between the artificially induced lead/lag and the DTW recovered lead/lag. For each parameter combination, 1,000 simulations are run. In panel (a): $\sigma_u = 1$, $\delta_1 = 5$, $\delta_2 = 10$, $W = 60$, $\psi = 6$, $N = 10,000$. In panel (b): $\delta_1 = 5$, $\delta_2 = 10$, $\sigma_{s_1} = 0.1$, $\sigma_{s_2} = 0.1$, $W = 60$, $\psi = 60$. In panel (c): $\delta_1 = 5$, $\delta_2 = 10$, $\sigma_{s_1} = 0.1$, $W = 60$, $\psi = 60$, $N = 10,000$.

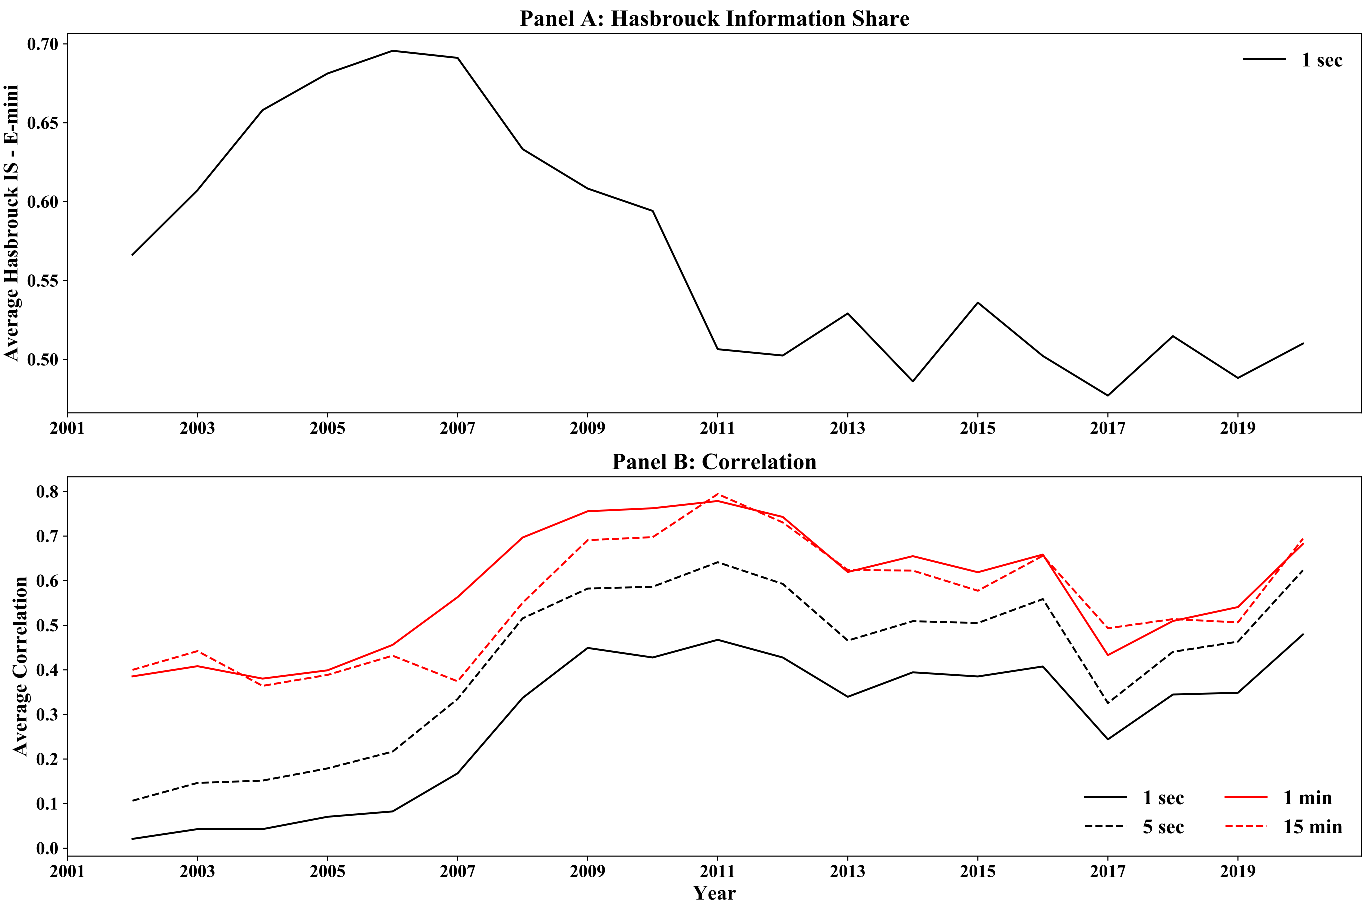

Figure 5: Hasbrouck information shares and correlations between E-mini S&P 500 futures and FTSE 100 futures

This figure presents the Hasbrouck Information Share (IS) and Pearson correlation between E-mini S&P 500 futures and FTSE 100 futures. Panel A presents the yearly average of the daily average of the upper and lower bounds of the Hasbrouck IS. For each day between January 1, 2002 and May 29, 2020, between 7:00 am UTC and 5:30 pm UTC, we use the one-second mid-quote log-returns to estimate the standard IS model. Panel B presents the yearly average of the daily correlation estimates between the E-mini futures and FTSE 100 futures at varying frequencies of mid-quote log-returns.

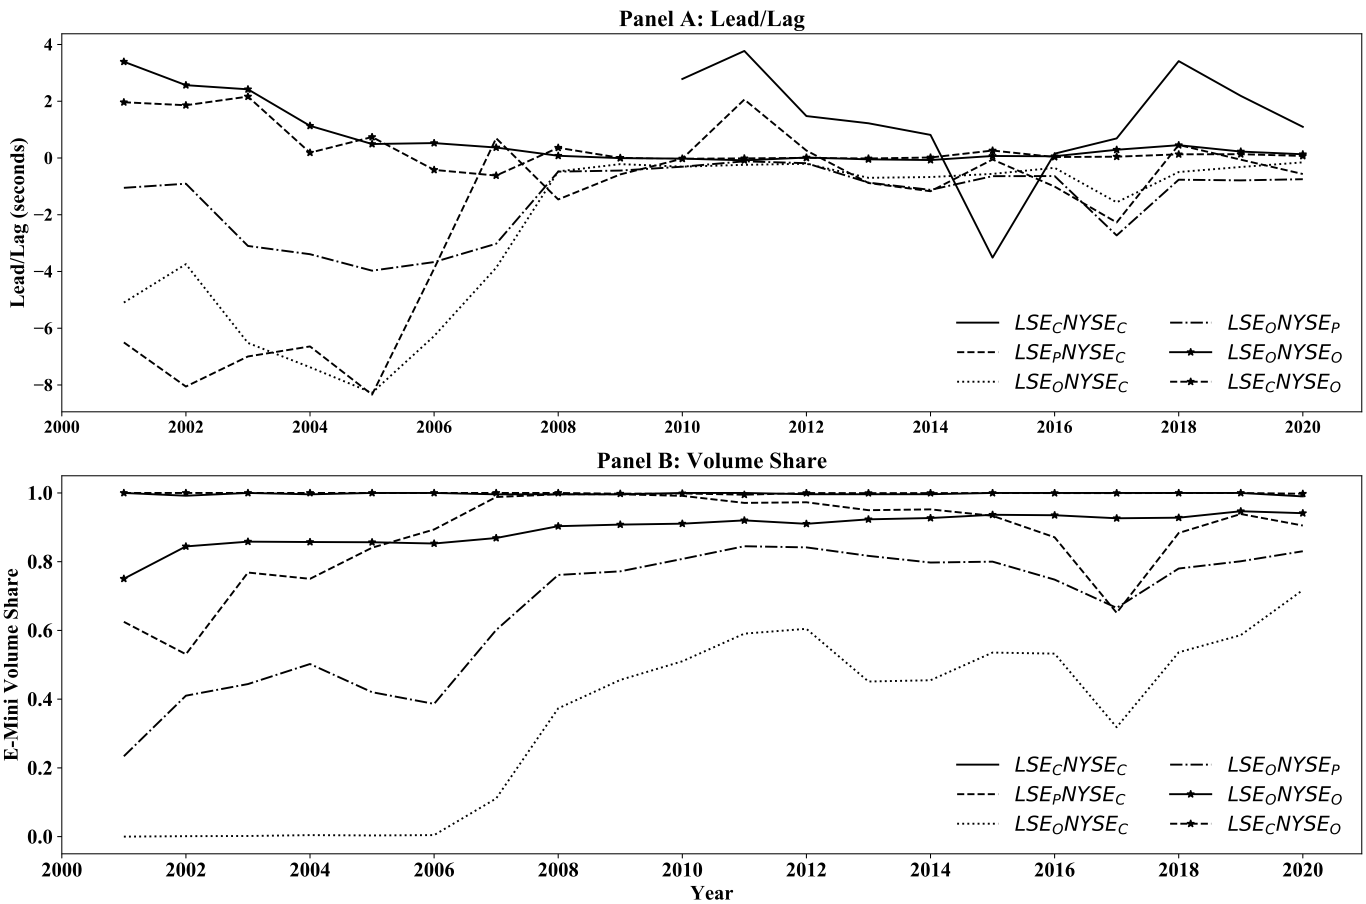

Figure 6: Yearly average of E-mini and FTSE 100 DTW estimated lead/lag and E-mini volume shares
This figure presents the yearly average of the daily median lead/lag between E-mini S&P 500 futures and FTSE 100 futures. For each day between November 13, 2001 and May 29, 2020, inclusive, intraday lead/lag values are estimating using one-second mid-quote log-returns on the E-mini and FTSE 100 futures, using DTW. The daily median of the one-second lead/lag is calculated using DTW and the yearly average is presented in Panel A. The E-mini volume share, presented in Panel B, is measured as the proportion of E-mini futures traded volume to the combined E-mini and FTSE 100 futures traded volume in each one-second interval. We split the trading day into six periods based on UTC time, and the current operating phase of the markets underlying the E-mini (New York Stock Exchange) and FTSE 100 (London Stock Exchange). These phases are pre-open (P), open (O), and closed (C).

Figure 7: Intraday lead/lag between E-mini and FTSE 100 - example A
This figure presents the cross-sectional average of the intraday one-second lead/lag between the E-mini S&P 500 futures and FTSE 100 futures, measured using DTW (solid lines). This figure also presents the cross-sectional average of the one-second intraday E-mini volume share (dotted lines). One-second mid-quote log-returns and volumes on the E-mini and FTSE 100 futures contracts are used. The black lines refer to the period 2014–2019, running from January 1, 2014 through May 31, 2019, using the months January through May in each year. The red lines refer to the period of January 1, 2020 through May 29, 2020. The shaded periods refer to the time at which either the London Stock Exchange or New York Stock Exchange are open. The left diagonal hatch is when the London Stock Exchange is open, the right diagonal hatch is when the New York Stock Exchange is open. We use Gaussian kernel smoothing to increase readability of the figure.

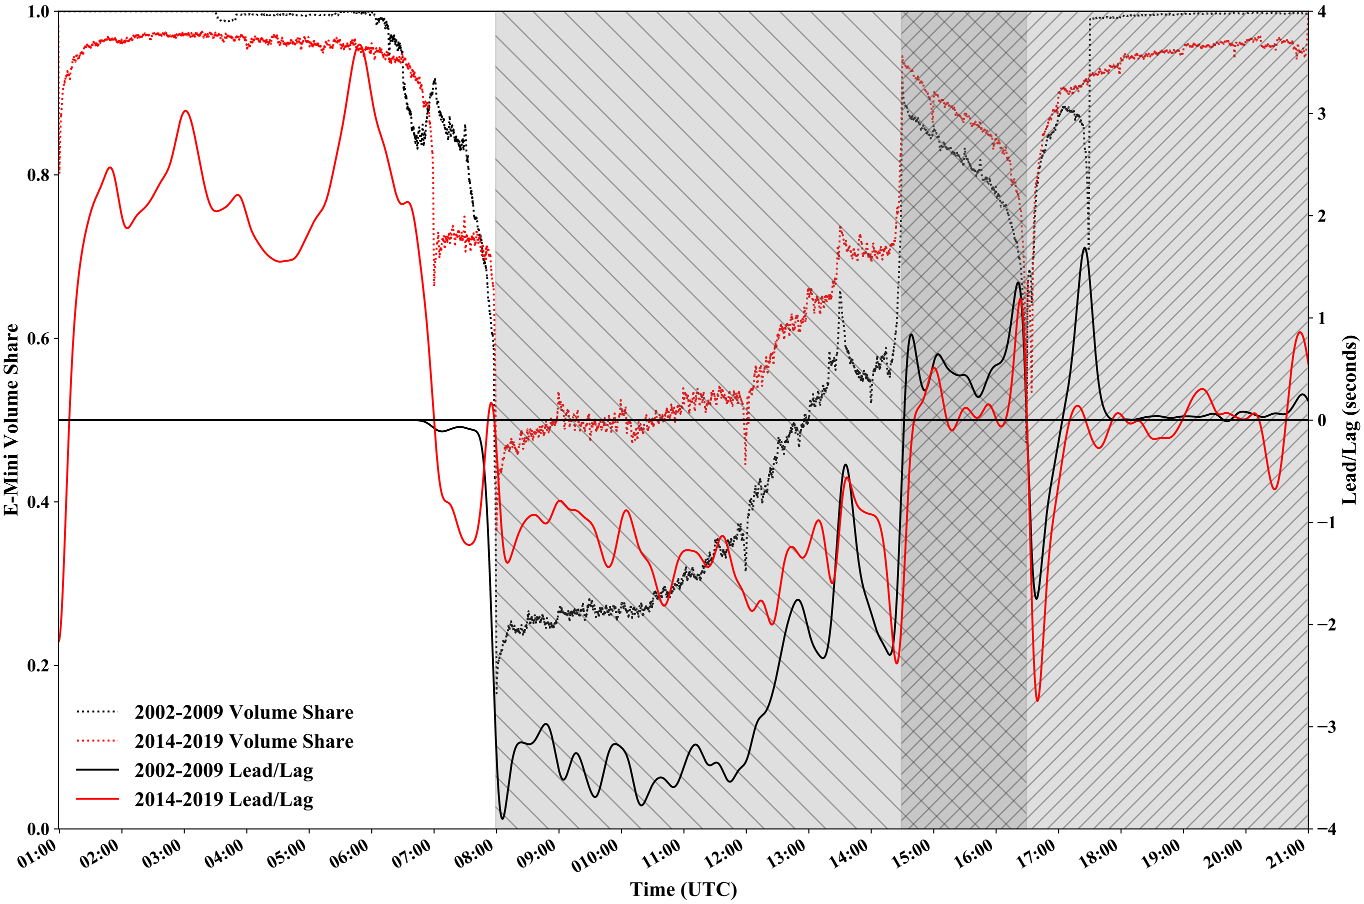

Figure 8: Intraday lead/lag between E-mini and FTSE 100 - example B
This figure presents the cross-sectional average of the one second intraday lead/lag between the E-mini S&P 500 futures and FTSE 100 futures, measured using DTW (solid lines). This figure also presents the cross-sectional average of the one-second intraday E-mini volume share (dotted lines). One-second mid-quote log-returns and volumes on the E-mini and FTSE 100 futures contracts are used. The black lines refer to the period from January 1, 2002 through December 31, 2009. The red lines refer to the period of January 1, 2014 through December 31, 2019. The shaded periods refer to the time at which the London Stock Exchange or New York Stock Exchange are open. The left diagonal hatch is when the London Stock Exchange is open, the right diagonal hatch is when the New York Stock Exchange is open. We use Gaussian kernel smoothing to increase readability of the figure.

Online Appendix

Table A.1: Adjusting for bias in the DTW algorithm

This table presents the results for a constant lead/lag simulation, in which we alternate the order of the two time-series in the calculation of the lead/lag using DTW. We simulate two assets which share a common fundamental value. We use the parameters $N = 10,000$, $u_t = 1$, $s_{1,t} = s_{2,t} = 0.5$ and $W = 60$. We induce a set of constant lead/lag between the two time-series in the range of -5 to +5, in increments of one unit. The second column contains the DTW lead/lag in which the order of the time-series in the DTW algorithm is time-series one first, and time-series two second. The third column contains the DTW lead/lag in which the order of the time-series is time-series two first, and time-series one second. The fourth column presents the difference of column two and column three divided by two.

True lead/lag	p_1/p_2 DTW lead/lag	p_2/p_1 DTW lead/lag	Bias-adjusted lead/lag
5	4.903	-5.095	4.999
4	3.903	-4.095	3.999
3	2.904	-3.095	2.999
2	1.905	-2.096	2.000
1	0.903	-1.094	0.999
0	-0.094	-0.097	0.002
-1	-1.095	0.904	-0.999
-2	-2.097	1.906	-2.001
-3	-3.097	2.906	-3.001
-4	-4.096	3.905	-4.000
-5	-5.096	4.904	-5.000

Table A.2: Portfolio characteristics of stocks sorted by DTWT

Each month from July 1927 to November 2019, two sets of decile portfolios are formed by sorting stocks based on DTWT over the past month. DTWT is the DTW t-statistic. This table presents the time-series average across the months of the median value of each characteristic within each decile portfolio.

Decile	β_{DTWT}	β_{DTW}	β_{DTWR}	β	β_{DIM}	β_{SW}	MAX	MIN	SIZE	ILLIQ	BM	MOM	REV	IVOL	TSKEW	SSKEW	PRC
Low β_{DTWT}	15.41	0.23	0.79	0.69	0.89	0.69	0.032	-0.031	10.17	4.37	0.73	-0.06	-0.96	0.36	0.63	-3.03	10.17
2	17.75	0.30	0.96	0.78	0.99	0.77	0.031	-0.030	10.44	1.47	0.70	-0.02	-0.43	0.36	0.46	-3.33	11.28
3	19.17	0.35	1.05	0.81	1.01	0.79	0.031	-0.029	10.57	1.01	0.71	0.01	-0.13	0.35	0.43	-3.40	11.95
4	20.33	0.38	1.11	0.83	1.02	0.81	0.031	-0.028	10.63	0.93	0.72	0.04	0.13	0.34	0.41	-3.44	12.77
5	21.42	0.41	1.15	0.84	1.03	0.82	0.030	-0.028	10.72	0.92	0.72	0.06	0.23	0.33	0.41	-3.46	13.99
6	22.49	0.42	1.17	0.86	1.04	0.83	0.030	-0.028	10.81	0.86	0.71	0.09	0.43	0.33	0.40	-3.43	15.35
7	23.61	0.42	1.18	0.87	1.05	0.84	0.030	-0.028	10.92	0.79	0.70	0.12	0.57	0.33	0.39	-3.54	16.79
8	24.87	0.41	1.18	0.89	1.06	0.86	0.030	-0.028	11.05	0.72	0.69	0.15	0.75	0.32	0.38	-3.51	18.70
9	26.44	0.40	1.16	0.91	1.08	0.87	0.030	-0.027	11.20	0.64	0.68	0.19	0.91	0.32	0.37	-3.57	20.81
High β_{DTWT}	29.01	0.36	1.12	0.93	1.09	0.89	0.030	-0.027	11.44	0.61	0.65	0.26	1.16	0.31	0.34	-3.46	24.51

Table A.3: Firm-level cross-sectional regressions: DTWT

This table presents the average coefficient estimates from monthly Fama and MacBeth (1973) regressions. Each month from July 1926 to December 2019 (regressions using Book-to-market is from July 1962 to December 2019), we regress excess stock returns during that month on lagged DTW $t-$statistic and several control variables defined in 5. Each row in the table reports the time-series averages of the cross-sectional regression slope coefficients and their associated Newey-West (1987) t-statistics, adjusting using 6 lags, in parentheses. The R-squared value for each regression is reported in the far right column. ***, ** and * indicate statistical significance at 1%, 5% and 10% levels, respectively.

	β_{DTWT}	SIZE	ILLIQ	MOM	IVOL	REV	MAX	BM	MIN	ISKEW	TSKEW	SSKEW	β	β_{DTW}	R^2
(i)	0.040														0.89%
	(3.57)***														
(ii)	0.063	-0.204													2.67%
	(5.91)***	(-4.32)													
(iii)	0.054	-0.085	0.047												4.31%
	(5.29)***	(-2.22)	(3.70)												
(iv)	0.036	-0.093	0.052	1.160											5.42%
	(4.73)***	(-2.67)	(4.29)	(5.38)											
(v)	0.030	-0.200	0.087	1.047	-2.318										6.47%
	(4.00)***	(-6.80)	(6.94)	(4.86)	(-11.13)										
(vi)	0.027	-0.144	0.083	0.693	-1.804	-0.058									7.42%
	(4.35)***	(-5.09)	(6.42)	(3.10)	(-7.57)	(-13.93)									
(vii)	0.030	-0.157	0.114	0.641	-1.392	-0.040	-10.696								5.79%
	(3.89)***	(-4.75)	(5.78)	(3.69)	(-4.38)	(-6.65)	(-2.04)								
(viii)	0.027	-0.130	0.129	0.504	-1.326	-0.048	-8.134	0.059							6.70%
	(3.55)***	(-3.88)	(4.71)	(2.67)	(-3.60)	(-6.81)	(-1.41)	(0.75)							
(ix)	0.027	-0.131	0.130	0.521	-1.124	-0.052	-7.034	0.056	4.251						6.92%
	(3.69)***	(-3.91)	(4.61)	(2.82)	(-2.45)	(-7.69)	(-1.32)	(0.72)	(0.90)						
(x)	0.026	-0.137	0.134	0.575	-1.091	-0.053	-5.924	0.058	5.748	-0.055					7.05%
	(3.38)***	(-4.11)	(4.63)	(3.08)	(-2.39)	(-7.88)	(-1.11)	(0.76)	(1.18)	(-2.37)					
(xi)	0.025	-0.139	0.136	0.589	-1.085	-0.054	-5.429	0.057	6.273	0.089	-0.166				7.14%
	(3.23)***	(-4.18)	(4.62)	(3.16)	(-2.39)	(-7.96)	(-1.02)	(0.74)	(1.28)	(1.16)	(-2.10)				
(xii)	0.025	-0.140	0.123	0.570	-1.129	-0.054	-4.989	0.060	6.052	0.122	-0.200	-0.005			7.33%
	(3.27)***	(-4.27)	(4.63)	(2.98)	(-2.47)	(-7.80)	(-0.93)	(0.79)	(1.24)	(1.64)	(-2.52)	(-1.08)			
(xiii)	0.028	-0.148	0.124	0.614	-0.974	-0.055	-5.822	0.057	8.076	0.117	-0.200	-0.006	0.131		8.16%
	(3.72)***	(-4.12)	(4.79)	(3.68)	(-2.54)	(-8.13)	(-1.27)	(0.80)	(1.99)	(1.61)	(-2.56)	(-1.34)	(1.10)		
(xiv)	0.028	-0.146	0.121	0.639	-0.976	-0.055	-5.969	0.057	8.059	0.117	-0.202	-0.006	0.090	0.165	8.28%
	(3.76)***	(-4.12)	(4.74)	(3.95)	(-2.50)	(-8.23)	(-1.29)	(0.80)	(2.00)	(1.60)	(-2.58)	(-1.36)	(0.78)	(2.67)***	

Table A.4: Mis-estimation of betas

This table presents the difference in beta estimations between the standard beta measure and Dimson, Scholes-Williams and DTW beta estimations. Each month we divide the stock universe into small, mid large stocks based on 30% lower and 70% upper cutoff values using the market capitalization of stocks on the NYSE. We then calculate value-weighted percentage differences between each of three beta estimation methods and the standard beta within each portfolio. We first create a Full Premia factor, which takes the difference in return between the LH+HH portfolios and the HL+LL portfolios. We then create a return set for the small universe where we take the difference of the LH and the LL portfolios. We repeat the same for large stocks. We repeat this process for each of the DTW beta, standard beta, Dimson beta and Scholes-Williams beta. We then take the monthly difference in returns between the DTW factors and the other beta factors. Each row in the table reports the time-series average of the difference in return between the created risk factors and their associated Newey-West (1987) t-statistics, adjusting using 6 lags, in parentheses. The R-squared value for each regression is reported in the far right column. ***, ** and * indicate statistical significance at 1%, 5% and 10% levels, respectively. Values are in percent deviation from the standard beta value.

1927–1975	β_{DIM}	β_{SW}	β_{DTW}
Small	0.329	-0.078	0.739
	(7.61)***	(2.87)***	(9.25)***
Mid	0.245	-0.109	0.324
	(2.71)***	(1.45)	(2.88)***
Large	0.050	-0.196	-0.116
	(4.32)***	(-6.90)***	(-0.29)
1975–2019			
Small	1.042	0.252	1.569
	(3.78)***	(8.37)***	(6.98)***
Mid	0.340	0.061	0.175
	(7.46)***	(5.50)***	(4.04)***
Large	-0.046	-0.130	-0.218
	(-0.14)	(-0.69)	(-6.01)***
1927–2019			
Small	0.536	0.126	0.952
	(4.88)***	(7.78)***	(8.95)***
Mid	0.264	0.000	0.278
	(4.37)***	(1.93)*	(3.89)***
Large	0.004	-0.174	-0.155
	(2.26)***	(-3.58)***	(-3.13)***