Investment-Horizon Spillovers*

Alex Chinco† and Mao Ye‡

November 6, 2017

Abstract

This paper uses wavelets to decompose each stock's trading-volume variance into frequency-specific components. We find that stocks dominated by short-run fluctuations in trading volume have abnormal returns that are 1% per month higher than otherwise similar stocks where short-run fluctuations in volume are less important—i.e., stocks with less of a short-run tilt. And, we document that a stock's short-run tilt can change rapidly from month to month, suggesting that these abnormal returns are not due to some persistent firm characteristic that's simultaneously adding both short-run fluctuations and long-term risk.

JEL Classification: C55, C58, G12, G14

Keywords: Investment Horizon, Trading Volume, Wavelet Variance

*We thank John Campbell, Bruce Carlin, Ian Dew-Becker, Andrew Ellul, Fangjian Fu, Xavier Gabaix, Joel Hasbrouck, Ohad Kadan, Andrew Karolyi, Robert Korajczyk, Pete Kyle, Lubos Pastor, Ronnie Sadka, Allan Timmermann, Jeff Wurgler, Amir Yaron, and Haoxiang Zhu as well as seminar participants at Illinois, the NFA conference, the FRA conference, Washington University of St. Louis, Northwestern, and Michigan State for extremely helpful comments and suggestions. Tao Feng, Rukai Lou, Xin Wang, Robbie Xu, Fan Yang, Miles Zheng, and Chao Zi provided excellent research assistance. This research is supported by National Science Foundation grant 1352936 (joint with the Office of Financial Research at the U.S. Treasury Department). This work also uses the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number OCI-1053575. We thank David O'Neal of the Pittsburgh Supercomputer Center for his assistance with supercomputing, which was made possible through the XSEDE Extended Collaborative Support Service (ECSS) program.

Current Version: http://www.alexchinco.com/investment-horizon-spillovers.pdf

†UIUC College of Business. www.alexchinco.com; alexchinco@gmail.com; (916) 709-9934.

‡UIUC College of Business and NBER. www.yemaofin.com; maoye@illinois.edu; (217) 244-0474.

1 Introduction

Long-term fundamentals get revealed in quarterly earnings reports and annual product releases, but short-run traders can turn over "half their net holdings in 137 seconds (Kirilenko et al., 2011)." These two frequencies differ by 5 orders of magnitude:

$$\frac{1 \text{ trade}}{137 \text{ sec}} \times \frac{60 \text{ sec}}{1 \text{ min}} \times \frac{60 \text{ min}}{1 \text{ hr}} \times \frac{24 \text{ hr}}{1 \text{ day}} \times \frac{365 \text{ days}}{1 \text{ year}} \times \frac{1 \text{ year}}{1 \text{ ancmt}} \approx 2.3 \times 10^5$$

To put this difference in perspective, there have only been around 2×10^5 generations since the human/chimpanzee divergence (MacKay, 2003). So, long-term fundamentals quite literally evolve at the speed of speciation from short-run traders' point of view. And, because these frequencies are so different, some people now believe that short-run trading activity "doesn't affect [long-term] investors very much."[1] But, is this correct? Can investors at the 1-month horizon really ignore trading activity at the 1-minute horizon?

Motivated by this question, we study how short-run trading activity affects long-term returns by decomposing each stock's trading-volume variance into frequency-specific components using the wavelet-variance estimator (Percival, 1995). After introducing this new statistical tool, we present two main results. First, we show that stocks dominated by short-run fluctuations in trading volume have abnormal returns that are 1% per month higher than otherwise similar stocks where short-run fluctuations in volume are less important—i.e., stocks with less of a "short-run tilt". Second, we show that a stock's short-run tilt can change rapidly from month to month, suggesting these abnormal returns are not due to some persistent firm characteristic that's adding both short-run fluctuations and long-term risk.

New Estimator. What share of a stock's trading-volume variance comes from changes in volume from one minute to the next? What about from one hour to the next? Or, from one day to the next? We begin our analysis by introducing the wavelet-variance estimator to answer these questions. We use wavelets rather than sine and cosine waves to perform this frequency decomposition because Fourier methods require trading-volume variance to be constant, but modern equity markets are dominated by sporadic clusters of intense trading activity followed by long quiet periods (Easley et al., 2012). And, wavelets are the natural way to accommodate non-stationary trading-volume data.

This non-stationarity in trading volume is not just an empirical phenomenon, either. It is present in any model where informed traders receive "private information only at the beginning of the trading session (Vayanos, 2001, pg. 133)"—i.e., in any model where informed traders have a well-defined investment horizon, where there is a beginning and an end to the time interval when traders can exploit a piece of information. And, this fact motivates the

[1] Asness, C. and Mendelson, M. *High-Frequency Hyperbole.* Wall Street Journal. Apr 1, 2014.

link between frequency-specific fluctuations in volume and investment horizon.

For example, in Kyle (1985) an informed trader will choose his demand each period so that the marginal benefit of acquiring an additional share is exactly offset by the expected cost of the resulting price impact summed over all future trading periods. So, when there are fewer future trading periods to bear the costs of any price impact, the informed trader will submit more aggressive market orders, making extreme trading-volume realization more likely during the second half of his investment horizon. Thus, if a large fraction of trading-volume variance comes from changes in volume from one minute to the next, then there is likely a large fraction of informed traders with a 2-minute investment horizon. Whereas, if a larger fraction of variance comes from hour-to-hour or day-to-day changes in volume, then there's likely a larger fraction of informed traders with a 2-hour or 2-day investment horizon.

Descriptive Statistics. After introducing the wavelet-variance estimator, we next summarize how the distribution of investment horizons varies over the entire cross-section of stocks. These descriptive statistics about the distribution of investment horizons for the entire cross-section are new and interesting because learning about the distribution of traders' investment horizons used to mean finding data on their portfolio holdings, and this data is hard to come by. So, earlier pioneering work was confined to studying either a small collection of stocks or a narrow range of horizons. For example, Brogaard et al. (2014) studies the fraction of volume coming from high-frequency traders but for only 60 stocks. Alternatively, Lan et al. (2015) studies institutional-investor turnover but only at the quarterly horizon.

For each NYSE-listed stock from January 2002 to December 2010, we use the wavelet-variance estimator to compute the share of trading-volume variance in month m due to fluctuations at each frequency ranging from a minute to a day. Then, for each stock-month, we summarize the information in this distribution with a single slope coefficient, which we refer to as a stock's "short-run tilt". Stocks with a larger short-run tilt in month m have more of their trading-volume variance concentrated at higher frequencies. And, when we sort stocks into quintiles based on their short-run tilt, we find that the typical stock in the highest quintile is traded 5-times faster than the typical stock in the lowest quintile. This difference is well outside the 95% confidence interval for what we should expect due to measurement error, and it persists when we control for firm characteristics, liquidity, and news announcements.

Horizon Spillover. Finally, we present our two main results. First, we show that the quintile of stocks with the largest short-run tilt has abnormal returns that are 1% per month higher than the quintile of stocks with the smallest short-run tilt. These abnormal returns are not explained by firm characteristics, news announcements, or standard liquidity measures.

But, although standard liquidity measures are computed at the daily horizon (Amihud, 2002; Pastor and Stambaugh, 2003; Sadka, 2006), what a monthly investor really cares about is a stock's liquidity at the exact time he decides to place an order. So, we use data from Ancerno Ltd. to measure the standard deviation of the difference between a stock's quoted price immediately prior an order and the price at which this order gets executed. We find that this more-granular liquidity measure explains around half of the abnormal returns associated with short-run tilt. This suggests one plausible way that trading activity at the 1-minute timescale might spillover and affect monthly returns.

Second, we document that the names of the stocks with the largest short-run tilt change rapidly from month to month. There is only a 32% chance that a stock belonging to the quintile with the largest short-run tilt in month m still belongs to this quintile in month $\{m+1\}$. We view this rapid churn as a new stylized fact worthy of further investigation. It is also a stylized fact supporting the hypothesis that the high abnormal returns earned by stocks with the largest short-run tilt are not due to some persistent firm characteristic that's simultaneously adding both short-run fluctuations and long-term risk.

1.1 Literature

This paper is related to several strands of literature.

Investment Horizon. First, this paper connects to work on the asset-pricing implications of heterogeneous investment horizons. On the empirical side, papers like Jagannathan and Wang (2007), Beber et al. (2012), van Binsbergen et al. (2012), Cella et al. (2013), Kamara et al. (2015), and Dew-Becker and Giglio (2016) evaluate the cross-horizon implications of popular asset-pricing models. On the theoretical side, Crouzet et al. (2017) highlights the distinguishing features of a model with multi-frequency trade.

We add to this literature by offering a new proxy for the distribution of traders' investment horizons. The ability to measure the entire distribution of investment horizons for the entire cross-section of stocks has a wide range of potential applications. For instance, one of the key constraints that arbitrageurs face in the limits-to-arbitrage models is a short investment horizon (De Long et al., 1990). There is also literature looking at the optimal response of firm managers to shareholders with biased short-term beliefs (Baker and Wurgler, 2013). And, this rational-manager vs biased-shareholder paradigm fits into an even larger literature looking at feedback effects between financial markets and the real economy (Chen et al., 2007) where the decisions of myopic investors also play a key role.

Trading Volume. In addition, this paper is related to work linking trading volume and

stock returns. This literature dates back to Gallant et al. (1992), Campbell et al. (1993), Andersen (1996), and Lo and Wang (2000). While related, the current paper also differs from this earlier work in an important way. These earlier papers focus on the link between trading volume and stock returns at the same horizon. For example, Lo and Wang (2000) looks at weekly turnover and weekly returns. By contrast, the current paper studies the effect of trading activity at shorter timescales on returns at longer timescales.

Wavelet Variance. Finally, this paper belongs to a growing list of papers that use wavelets to answer asset-pricing questions. For example, see Gençay et al. (2001), which analyzes a variety of financial applications for wavelets, Hasbrouck (2015), which uses wavelets to analyze high-frequency patterns in the limit-order book, and Bandi et al. (2015), which uses wavelets to forecast stock returns. This paper is the first to apply the wavelet-variance estimator to analyze trading-volume data. This is an application for which the wavelet-variance estimator is uniquely suited because trading-volume variance is time varying.

1.2 Notation

We use the following notation throughout the paper. The Gr$\epsilon\epsilon\kappa$ alphabet denotes parameters and the Roman alphabet denotes variables. $m = 1, \ldots, M$ counts months while $t = 1, \ldots, T$ counts minutes within each month. Monthly variables are written in `teletype font`. A $\widehat{\text{hat}}$ denotes estimated parameters, a $\overline{\text{bar}}$ denotes sample averages, and a $\widetilde{\text{tilde}}$ denotes variables normalized to have zero mean and unit variance.

2 New Estimator

What share of a stock's trading-volume variance comes from changes in volume from one minute to the next? What about from changes in volume from one hour to the next? Or, from one day to the next? In this section, we motivate why answering these questions might be informative about the distribution of traders' investment horizons. Then, we explain how the wavelet-variance estimator decomposes a stock's trading-volume variance into frequency-specific components and why we use wavelets rather than waves in our analysis.

2.1 Frequency and Horizon

A trader's investment horizon is the length of time from when he learns a new piece of information until when this information becomes public. And, as Vayanos (2001) points out, trading volume will be non-stationary in any model where traders have a well-defined

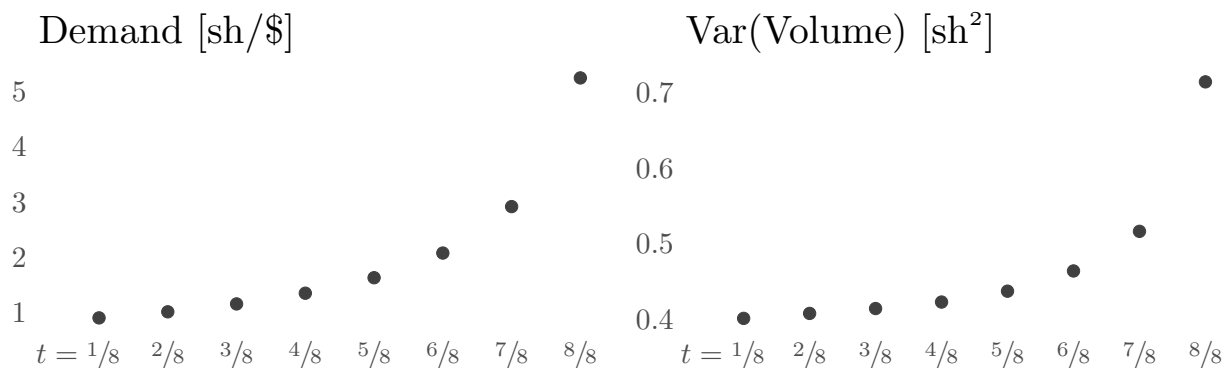

Figure 1: Informed Demand. *Number of shares that an informed trader will demand per dollar of fundamental value as a function of t in a Kyle (1985) model where $T = 8$. Reads: "When there is less time remaining, the informed trader submits more aggressive demand orders."* **Trading-Volume Variance.** *Realized trading-volume variance as a function of t in 10,000 simulations of the same Kyle (1985) model where $T = 8$. Reads: "Trading-volume variance increases when there is less time remaining because the informed trader submits more aggressive demand orders."*

investment horizon. We now outline why this observation suggests that frequency-specific fluctuations in trading volume will be informative about traders' investment horizons.

Fluctuations in Volume. Consider the standard discrete-time Kyle (1985) model with a single stock and time indexed by $t \in \{1/T, 2/T, \ldots, T/T\}$. Investment horizon in this model is captured by the units on time. If $t = T/T = 1$ occurs one hour from now, then informed traders in the model have an investment horizon of an hour. Whereas, if $t = 1$ occurs one day from now, then they have an investment horizon of a day.

The informed trader in this model will choose his demand each period so that the marginal benefit of acquiring an additional share in the current period is exactly offset by the expected cost of the resulting price impact summed over all current and future trading periods. So, when choosing his demand in the first period $t = 1/T$, the informed trader knows that he will bear the cost of any price impact in all N trading periods. And, when choosing his demand in period $t = 1$, he knows that he will bear the cost only in the current trading period. Thus, later in his investment horizon, the informed trader will submit more aggressive market orders since the cost of the resulting price impact will be felt in fewer periods (left panel of Figure 1), which will lead to an upward trend in trading-volume variance (right panel).

Frequency Decomposition. If trading-volume variance is trending up, then there will tend to be more extreme trading-volume realizations in the second half of the informed trader's investment horizon. So, the canonical Kyle (1985) model suggests a link between frequency-specific fluctuations in trading volume and traders' investment horizons. If there's a lot of variance coming from changes in volume from one 30-minute interval to the next, then

$t = 1$ likely corresponds to one hour for lots of informed traders. Whereas, if there's a lot of variance coming from changes in volume from one 12-hour interval to the next, then $t = 1$ likely corresponds to one day for lots of informed traders.

2.2 Spectral Variance

But, how should we measure the share of trading-volume variance coming from fluctuations at different frequencies? At first, it might seem natural to use Fourier methods. But, we now show why the time-varying nature of trading-volume variance makes this is a bad idea.

Variance Decomposition. Here is how to decompose trading-volume variance into frequency-specific components using Fourier methods. Suppose we have $T \geq 2$ minutes of trading-volume data for a single stock where $T = 2^F$ for some $F \in \{1, 2, \ldots\}$. The discrete Fourier transform, ν_f, of this demeaned trading-volume time series at frequency f is defined as:

$$\nu_f \stackrel{\text{def}}{=} \frac{1}{\sqrt{T}} \cdot \sum_{t=1}^{T/2} \{vlm(t) - \overline{vlm}\} \cdot e^{-2 \cdot \pi \cdot i \cdot f \cdot t/T}$$

Parseval's theorem relates the variation in trading volume over time periods to the variation in volume across frequencies (Mallat, 1999, pg. 39):

$$\text{Var}(vlm) = \frac{1}{T} \cdot \sum_{f=1}^{T/2} |\nu_f|^2$$

And, we can estimate $|\nu_f|^2$ for each $f = \{1, \ldots, T/2\}$ by regressing demeaned trading volume on scaled sine and cosine waves, $s_f(t) \stackrel{\text{def}}{=} \sin(2 \cdot \pi \cdot f \cdot t/T)$ and $c_f(t) \stackrel{\text{def}}{=} \cos(2 \cdot \pi \cdot f \cdot t/T)$:

$$vlm(t) - \overline{vlm} = \sum_{f=1}^{T/2} \left\{ \hat{\alpha}_f \cdot s_f(t) + \hat{\beta}_f \cdot c_f(t) \right\} \tag{1}$$

So, by Euler's identity, $e^{i \cdot x} = \cos(x) + i \cdot \sin(x)$, we can compute the amount of trading-volume variance at frequency f as:

$$\text{SVar}_f(vlm) \stackrel{\text{def}}{=} \tfrac{1}{T} \cdot \{\hat{\alpha}_f^2 + \hat{\beta}_f^2\} \tag{2}$$

This approach to decomposing variance into frequency-specific components is known as the spectral-variance estimator. And, we define the corresponding estimate of the fraction of trading-volume variance at frequency f as:

$$\text{SFrac}_f(vlm) \stackrel{\text{def}}{=} \frac{\text{SVar}_f(vlm)}{\sum_{f'} \text{SVar}_{f'}(vlm)} \tag{3}$$

Unstable Results. Unfortunately, because sine and cosine waves fluctuate in the exact same way over the course of the entire sample period, the spectral-variance estimator is

unstable when trading-volume variance moves around over time. And, we know that equity markets are dominated by periods of intense trading activity followed by long quiet periods with little action (Easley et al., 2012). We just saw that informed traders will submit more aggressive demand orders later in the lifecycle of a single piece of private information (left panel of Figure 1). What's more, the unexpected news that constitutes an informed trader's private information cannot arrive at regular intervals by definition. So, there will also be irregular gaps between the frequency-specific fluctuations in trading volume documented in the right panel of Figure 1. These facts manifest as evidence of regime switching (Epps and Epps, 1976; Harris, 1987; Andersen, 1996) and as failures to reject non-stationary (Gallant et al., 1992; Campbell et al., 1993; Lo and Wang, 2000).

2.3 Wavelet Variance

However, if time-varying trading-volume variance is the problem, then maybe we can solve it by applying the spectral-variance estimator separately to short sub-samples of data where trading-volume variance is almost constant? This is the intuition behind our application of the wavelet-variance estimator (Percival, 2016).

General Idea. Using the wavelet-variance estimator involves regressing a stock's demeaned trading volume on frequency-specific variables, $w_{f,\ell}(t)$, just like in Equation 1:

$$vlm(t) - \overline{vlm} = \sum_{f=1}^{F} \sum_{\ell=1}^{2^{f-1}} \hat{\theta}_{f,\ell} \cdot w_{f,\ell}(t) \tag{4}$$

And, just like in Equation 2, when using the wavelet-variance estimator, each frequency's contribution to a stock's total trading-volume variance will be an average of squared coefficients, $\frac{1}{T} \cdot \sum_{\ell=1}^{2^{f-1}} \hat{\theta}^2_{f,\ell}$. The only difference between decomposing trading-volume variance with waves and wavelets is that, when using wavelets, the right-hand-side variables in Equation 1 have to be indexed by both their frequency and their location in time.

Wavelets are halfway between time fixed effects and sine waves. A time fixed effect, $1_{\{t=\ell\}}$, only has non-zero output when $t = \ell$; whereas, the sine function oscillates at frequency f for any time $t \in \mathcal{R}$. Wavelets split the difference between these two extremes, returning values that oscillate at frequency f only during a short time period indexed by ℓ (for "location" in time). And, because wavelets are confined to fluctuating only at a particular location in time, using them to decompose trading-volume variance is like applying the spectral-variance estimator separately to short subsamples of data where the variance is almost constant.

Defining "Halfway". Let's now define the idea of a wavelet in more detail. Assume that the number of trading periods can be written as $T = 2^F$ for some highest frequency

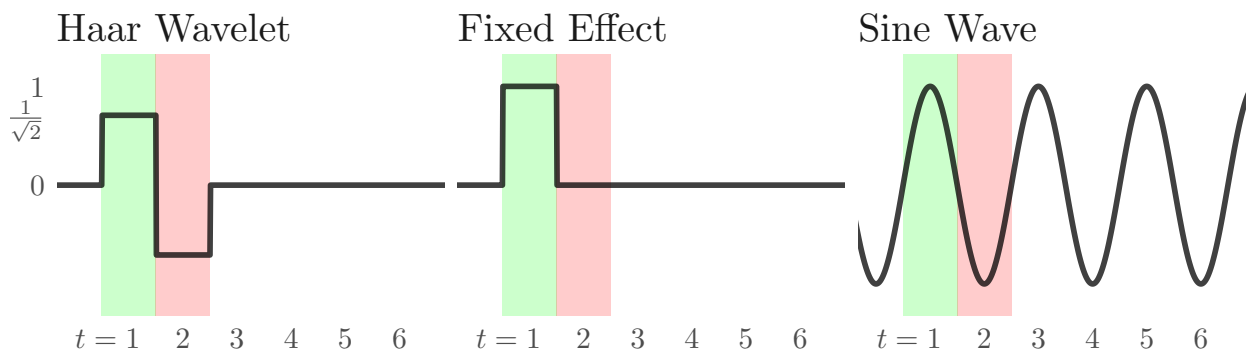

Figure 2: *A wavelet and some non-wavelets.* **Haar Wavelet.** *A variable that is* $1/\sqrt{2}$ *during the first period,* $-1/\sqrt{2}$ *during the second period, and* 0 *otherwise.* **Fixed Effect.** *A variable that is* 1 *during the first period and* 0 *otherwise. It's not a wavelet because it is not mean zero.* **Sine Wave.** $\sin(2 \cdot \pi \cdot t)$. *It's not a wavelet because it is not normalizable.*

$F \in \{1, 2, \ldots\}$. A wavelet is a function, $w_{f,\ell}(t)$, defined at each time $t \in \{1, \ldots, T\}$ with both a characteristic frequency, $f \in \{1, \ldots, F\}$, and a characteristic location in time, $\ell(f) = \{1, \ldots, 2^{f-1}\}$, that satisfies the following three properties.

First, a wavelet has to be orthogonal:

$$0 = \sum_{t=1}^{T} w_{f,\ell}(t) \cdot w_{f',\ell'}(t) \quad \text{whenever} \quad f \neq f' \text{ or } \ell \neq \ell' \tag{P1}$$

Orthogonality implies that each of the wavelets used as right-hand-side variables in Equation 4 will capture different information about fluctuations in trading volume. Time fixed effects, such as the one in the center panel of Figure 2, also satisfy this orthogonality condition.

Second, a wavelet has to be mean-zero:

$$0 = \sum_{t=1}^{T} w_{f,\ell}(t) \tag{P2}$$

This is where the "wave" in "wavelet" comes from. A wavelet has to peak above zero and then dip below zero, and a wavelet's f index determines the length of this cycle. The time fixed effect in Figure 2 does not satisfy this second property.

Third, a wavelet has to be normalizable:

$$1 = \sum_{t=1}^{T} w_{f,\ell}(t)^2 \quad \text{for any} \quad T \geq 2^f \tag{P3}$$

This is where the diminutive "let" in "wavelet" comes from. To satisfy this condition, a function has to be localized in time, and a wavelet's ℓ index determines this location. The canonical sine wave is not localized in time and so is not normalizable. As the sample size gets larger and larger, a sine wave's quadratic variation continues to grow and grow.

There are many functions that satisfy these three properties. We are going to use Haar wavelets because we want to emphasize the fact that you can think about wavelets as a special kind of right-hand-side variable that's a cross between a fixed effect and sine wave.

Definition 1 (Haar Wavelet)**.** *A Haar wavelet is a piecewise-constant function of time defined on a time series with $T = 2^F$ observations:*

$$w_{f,\ell}(t) \stackrel{\text{def}}{=} \begin{cases} \sqrt{2^{f-1}/T} & \textit{if} \quad \ell - 1 \ < t/T \cdot 2^{f-1} \leq \ell - 1/2 \\ -\sqrt{2^{f-1}/T} & \textit{if} \quad \ell - 1/2 < t/T \cdot 2^{f-1} \leq \ell \\ 0 & \textit{else} \end{cases}$$

The wavelet's frequency index can take on values $f \in \{1, \ldots, F\}$, and at a given frequency the wavelet's location index can take on values $\ell(f) \in \{1, \ldots, 2^{f-1}\}$.

Variance Decomposition. We now use these three wavelet properties to derive a frequency-specific variance decomposition analogous to Equation 2 but using wavelets instead of sine and cosine waves. We can express a stock's demeaned trading volume in a wavelet basis à la Equation 4, which means that we can rewrite a stock's trading-volume variance as follows:

$$\text{Var}(vlm) = \text{Var}\left(\sum_{f=1}^{F} \sum_{\ell=1}^{2^{f-1}} \hat{\theta}_{f,\ell} \cdot w_{f,\ell} \right)$$

The fact that wavelets are orthogonal to one another (P1) then implies that we can rewrite this expression as the sum of variances rather than the variance of sums:

$$\sum_{f=1}^{F} \sum_{\ell=1}^{2^{f-1}} \hat{\theta}^2_{f,\ell} \cdot \text{Var}(w_{f,\ell})$$

And, the fact that wavelets are both mean-zero (P2) and normalizable (P3) further implies that $\text{Var}(w_{f,\ell}) = 1/T$. So, we can further simplify this expression as follows:

$$\sum_{f=1}^{F} \left\{ \frac{1}{T} \cdot \sum_{\ell=1}^{2^{f-1}} \hat{\theta}^2_{f,\ell} \right\}$$

In other words, a stock's trading-volume variance can be decomposed into F different frequency-specific components by looking at the average of squared wavelet coefficients at each frequency, with the component at frequency f representing the amount of trading-volume variance that can be explained by comparing successive periods of length $T/2^f$.

Definition 2 (Wavelet-Variance Estimator)**.** *The wavelet-variance estimator at frequency $f \in \{1, \ldots, F\}$ for a trading-volume time series with $T = 2^F$ observations is defined as:*

$$\text{WVar}_f(vlm) \stackrel{\text{def}}{=} \frac{1}{T} \cdot \sum_{\ell=1}^{2^{f-1}} \hat{\theta}^2_{f,\ell} \tag{5}$$

And, we define the corresponding fraction of trading-volume variance at frequency f as:

$$\text{WFrac}_f(vlm) \stackrel{\text{def}}{=} \frac{\text{WVar}_f(vlm)}{\sum_{f'} \text{WVar}_{f'}(vlm)} \tag{6}$$

Note that, in contrast to this spectral-variance estimator, the wavelet-variance estimator is well-defined even when trading-volume variance is time varying (Percival, 1995). If movements in trading-volume variance occur at frequency f_0, then trading-volume variance will be effectively constant at all higher frequencies $f > f_0$. So, the average of squared wavelet coefficients at frequency $f > f_0$ will be well defined and represent the average contribution of fluctuations at frequency f to changes in trading volume over all 2^{f-1} locations in time during the entire sample period (Serroukh et al., 2000).

2.4 An Example

To make this definition concrete, we now analyze the variance of a sample trading-volume time series with $T = 8$ observations.

Haar Wavelets. There are Haar wavelets at $\log_2(8) = 3$ different frequencies associated with an 8-period sample. At frequency $f = 1$ there is a single wavelet, $1 = 2^{1-1}$:

$$\begin{bmatrix} \underset{w_{1,1}(1)}{1/\sqrt{8}} & \underset{w_{1,1}(2)}{1/\sqrt{8}} & \underset{w_{1,1}(3)}{1/\sqrt{8}} & \underset{w_{1,1}(4)}{1/\sqrt{8}} & \underset{w_{1,1}(5)}{-1/\sqrt{8}} & \underset{w_{1,1}(6)}{-1/\sqrt{8}} & \underset{w_{1,1}(7)}{-1/\sqrt{8}} & \underset{w_{1,1}(8)}{-1/\sqrt{8}} \end{bmatrix}$$

This variable captures how the first half of the trading-volume time series differs from the second. At frequency $f = 2$ there are $2 = 2^{2-1}$ wavelets:

$$\begin{bmatrix} \underset{w_{2,1}(1)}{1/2} & \underset{w_{2,1}(2)}{1/2} & \underset{w_{2,1}(3)}{-1/2} & \underset{w_{2,1}(4)}{-1/2} & \underset{w_{2,1}(5)}{0} & \underset{w_{2,1}(6)}{0} & \underset{w_{2,1}(7)}{0} & \underset{w_{2,1}(8)}{0} \end{bmatrix}$$

$$\begin{bmatrix} \underset{w_{2,2}(1)}{0} & \underset{w_{2,2}(2)}{0} & \underset{w_{2,2}(3)}{0} & \underset{w_{2,2}(4)}{0} & \underset{w_{2,2}(5)}{1/2} & \underset{w_{2,2}(6)}{1/2} & \underset{w_{2,2}(7)}{-1/2} & \underset{w_{2,2}(8)}{-1/2} \end{bmatrix}$$

This pair of variables captures how the first quarter of the time series differs from the second as well as how the third quarter of the time series differs from the fourth. And finally, at frequency $f = 3$ there are $4 = 2^{3-1}$ wavelets:

$$\begin{bmatrix} \underset{w_{3,1}(1)}{1/\sqrt{2}} & \underset{w_{3,1}(2)}{-1/\sqrt{2}} & \underset{w_{3,1}(3)}{0} & \underset{w_{3,1}(4)}{0} & \underset{w_{3,1}(5)}{0} & \underset{w_{3,1}(6)}{0} & \underset{w_{3,1}(7)}{0} & \underset{w_{3,1}(8)}{0} \end{bmatrix}$$

$$\begin{bmatrix} \underset{w_{3,2}(1)}{0} & \underset{w_{3,2}(2)}{0} & \underset{w_{3,2}(3)}{1/\sqrt{2}} & \underset{w_{3,2}(4)}{-1/\sqrt{2}} & \underset{w_{3,2}(5)}{0} & \underset{w_{3,2}(6)}{0} & \underset{w_{3,2}(7)}{0} & \underset{w_{3,2}(8)}{0} \end{bmatrix}$$

$$\begin{bmatrix} \underset{w_{3,3}(1)}{0} & \underset{w_{3,3}(2)}{0} & \underset{w_{3,3}(3)}{0} & \underset{w_{3,3}(4)}{0} & \underset{w_{3,3}(5)}{1/\sqrt{2}} & \underset{w_{3,3}(6)}{-1/\sqrt{2}} & \underset{w_{3,3}(7)}{0} & \underset{w_{3,3}(8)}{0} \end{bmatrix}$$

$$\begin{bmatrix} \underset{w_{3,4}(1)}{0} & \underset{w_{3,4}(2)}{0} & \underset{w_{3,4}(3)}{0} & \underset{w_{3,4}(4)}{0} & \underset{w_{3,4}(5)}{0} & \underset{w_{3,4}(6)}{0} & \underset{w_{3,4}(7)}{1/\sqrt{2}} & \underset{w_{3,4}(8)}{-1/\sqrt{2}} \end{bmatrix}$$

These variables capture period-to-period differences—i.e., how the first period differs from the second, how the third period differs from the fourth, etc. . .

Variance Decomposition. Suppose we want to use these 7 wavelets to analyze a trading-

volume time series with most of its variance concentrated at the 1-period timescale:

$$\begin{bmatrix} \underset{vlm(1)}{200} & \underset{vlm(2)}{200} & \underset{vlm(3)}{200} & \underset{vlm(4)}{300} & \underset{vlm(5)}{200} & \underset{vlm(6)}{200} & \underset{vlm(7)}{200} & \underset{vlm(8)}{300} \end{bmatrix}$$

We say that this time series has most of its variation concentrated at the 1-period timescale because it is constant except for a +100 share shock in period $t = 4$ and a +100 share shock in period $t = 8$.

Regressing this demeaned trading-volume time series that on our 7 wavelets as in Equation 4 yields the following coefficients:

$$\begin{bmatrix} \underset{\hat{\theta}_{1,1}}{0} & \underset{\hat{\theta}_{2,1}}{-100/2} & \underset{\hat{\theta}_{2,2}}{-100/2} & \underset{\hat{\theta}_{3,1}}{0} & \underset{\hat{\theta}_{3,2}}{-100/\sqrt{2}} & \underset{\hat{\theta}_{3,3}}{0} & \underset{\hat{\theta}_{3,4}}{-100/\sqrt{2}} \end{bmatrix}$$

And, these coefficients have natural interpretations. For example, $\hat{\theta}_{1,1} = 0$ since on average trading volume during the first 4 minutes is the same as volume during the second 4 minutes. Likewise, $\hat{\theta}_{3,4} < 0$ since volume in the eighth period is greater than volume in the seventh.

The wavelet-variance estimator then implies that there is no trading-volume variance concentrated at frequency $f = 1$:

$$\text{Wvar}_1(vlm) = 1/8 \cdot \big\{0\big\}^2 = 0$$

It implies that some of the trading-volume variance is concentrated at frequency $f = 2$:

$$\text{Wvar}_2(vlm) = 1/8 \cdot \big\{-100/2\big\}^2 + 1/8 \cdot \big\{-100/2\big\}^2 = 625$$

This is the variance coming from successive 2-period increments having different amounts of volume. And, it implies that most of the trading-volume variance is concentrated at frequency $f = 3$ is:

$$\text{Wvar}_3(vlm) = 1/8 \cdot \big\{0\big\}^2 + 1/8 \cdot \big\{-100/\sqrt{2}\big\}^2 + 1/8 \cdot \big\{0\big\}^2 + 1/8 \cdot \big\{-100/\sqrt{2}\big\}^2 = 1250$$

This is the variance coming from successive periods having different amounts of volume.

We plot this trading-volume time series along with its associated wavelet-variance decomposition in the left panels of Figure 3. The multi-colored step function represents the Haar wavelets associated with the $2^{3-1} = 4$ different locations at frequency $f = 3$. Because each of these wavelets has a separate coefficient in Equation 4, slight changes in the timing of the spikes in volume do not affect the results. To see this, note that the fraction of trading-volume variance concentrated at each frequency is identical in the center panels of Figure 3 where the first +100 share spike in trading volume occurs at a slightly different time.

Spectral Instability. Figure 4 displays the analogous results using the spectral-variance estimator, which uses Fourier methods rather than wavelets. The right panel shows that both estimators give identical results when applied to data with constant variance. We

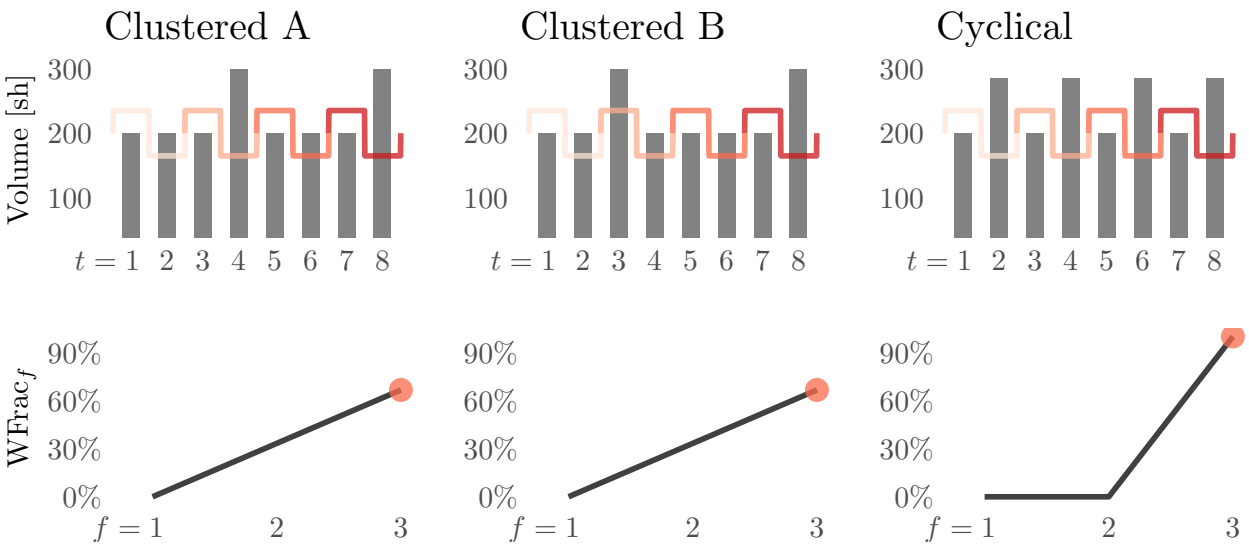

Figure 3: *Top panels depict* 3 *different trading-volume time series with* $T = 8$ *observations that are each dominated by fluctuations in volume from one minute to the next. Bottom panels show the fraction of trading-volume variance at each frequency as suggested by the wavelet-variance estimator defined in Equation 6. The multi-colored step-curve represents the* 4 *different wavelets at frequency* $f = 3$ *which capture fluctuations in trading-volume from one minute in Equation 4. The fraction of trading-volume variance at the 1-minute horizon is highlighted in orange in the bottom panel to make the comparison clear.* **Clustered A, B.** *A pair of time series with irregular spikes in trading volume. These time series correspond to a model where new information emerges at irregular intervals, like in the real world.* **Cyclical.** *A time series with cyclical spikes in trading volume every other minute. This time series corresponds to a model where new information emerges every* 120 *seconds like clockwork. Reads: "All three trading-volume time series have the majority of their variance concentrated at the highest frequency."*

know that variance is constant in the right panel because the time series has clear periodic structure. However, the left and center panels show that, when there are irregular spikes in trading volume, the output from the spectral-variance estimator is highly unstable because the spectral-variance estimator has no way of separately accounting for the contribution of fluctuations at frequency f during booms and lulls in trading volume. This fact is a direct consequence of the well-known uncertainty relation between representations of the same data in the time and frequency domains (Mallat, 1999, pg. 43).

2.5 Estimator Properties

We conclude this section by characterizing the wavelet-variance decompositions and associated confidence intervals of several commonly used data-generating processes. Later, these results will be useful for interpreting our empirical findings.

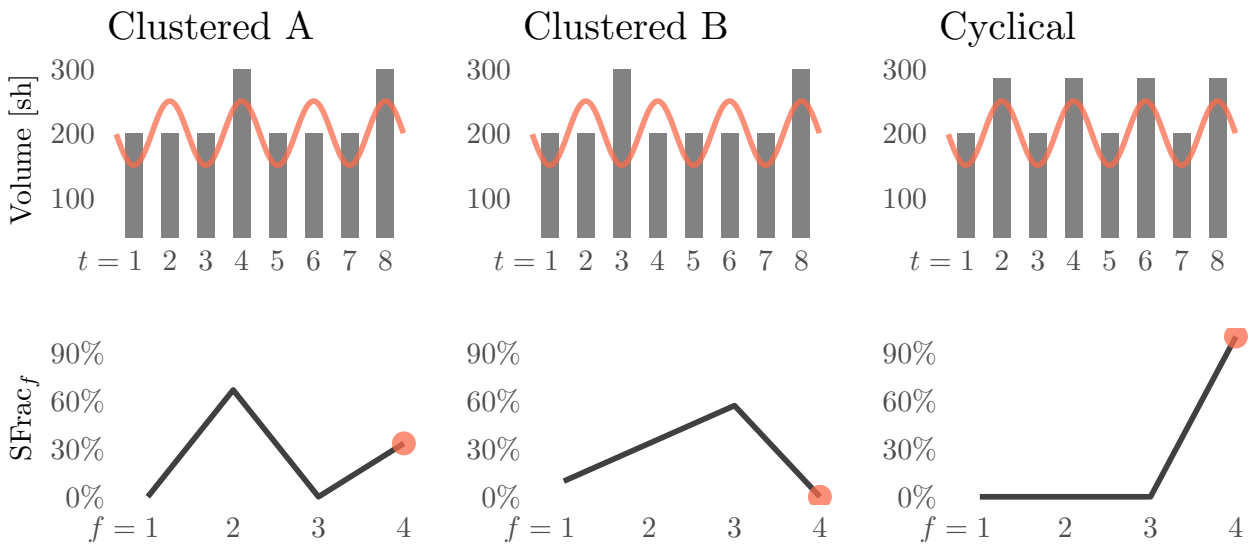

Figure 4: *Top panels depict the same* 3 *trading-volume time series as in Figure 3. Bottom panels show the fraction of trading-volume variance at each frequency as suggested by the spectral-variance estimator defined in Equation 3. The wavy orange line represents the scaled cosine function at frequency* $f = 4$ *which captures fluctuations in trading-volume from one minute to the next in Equation 1. The fraction of trading-volume variance at the 1-minute horizon is highlighted in orange in the bottom panel to make the comparison clear.* **Clustered A, B.** *A pair of time series with irregular spikes in trading volume. These time series correspond to a model where new information emerges at irregular intervals, like in the real world.* **Cyclical.** *A time series with cyclical spikes in trading volume every other minute. This time series corresponds to a model where new information emerges every* 120 *seconds like clockwork. Reads: "The spectral-variance estimator is highly unstable when there are irregular spikes in trading volume but works well when volume is cyclical."*

Special Cases. There are closed-form solutions for the wavelet-variance decompositions of several well-known stochastic processes (Gençay et al., 2001). For a white-noise process, the fraction of variance at frequency f is:

$$x(t) \overset{\text{iid}}{\sim} \mathrm{N}(0,1) \quad \Rightarrow \quad \mathrm{WFrac}_f(x) = 2^{f-1}/T \tag{7}$$

In other words, white noise is dominated by high-frequency fluctuations. Half of its variation will come from comparing successive periods. Any variation at lower frequencies will come from shocks in successive periods coincidentally lining up in the same direction, and the likelihood of longer and longer coincidences will decay at an exponential rate.

For a random walk, however, the fraction of variance concentrated at frequency f is:

$$\Delta x(t) \overset{\text{iid}}{\sim} \mathrm{N}(0,1) \quad \Rightarrow \quad \mathrm{WFrac}_f(x) = 1 - 2^{f-1}/T \tag{8}$$

A random walk represents the other extreme: it is a process dominated by low-frequency fluctuations. This makes sense because a random walk has a unit root, meaning that the effects of a shock at time t persist indefinitely. Half of the variation in a random walk will

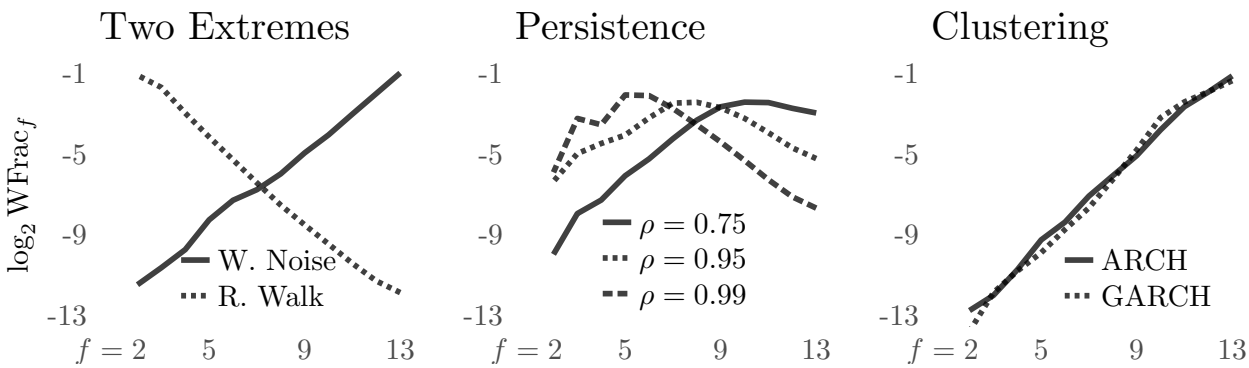

Figure 5: *Wavelet-variance decompositions of simulated data for commonly-used data-generating processes. Each simulated time series contains* $T = 2^{13}$*, and* $z(t) \overset{iid}{\sim} \mathrm{N}(0,1)$*.* **Two Extremes.** *Solid line is decomposition of a white-noise process,* $x(t) = z(t)$*. Dashed line is decomposition of a random walk,* $\Delta x(t) = z(t)$*. Reads: "White noise is dominated by high-frequency variation (slope* $= +1$*), and random walks are dominated by low-frequency variation (slope* $= -1$*)."* **Persistence.** *Decompositions of* AR(1) *processes with* $\rho \in \{0.75, 0.95, 0.99\}$*,* $x(t) = \rho \cdot x(t-1) + z(t)$*. Reads: "We can smoothly move between these two extremes by adjusting the persistence parameter,* $\rho \in [0,1]$*, with larger slope implying lower persistence."* **Clustering.** *Solid line is decomposition of an* ARCH(1) *process,* $x(t) = \sigma(t) \cdot z(t)$ *and* $\sigma(t)^2 = 0.25 + 0.90 \cdot z(t-1)^2$*. Dashed line is decomposition of a* GARCH(1,1) *process,* $x(t) = \sigma(t) \cdot z(t)$ $\sigma(t)^2 = 0.25 + 0.45 \cdot \sigma(t-1)^2 + 0.45 \cdot z(t-1)^2$*. Reads: "The wavelet-variance estimator is robust to allowing for time-varying variance."*

come from comparing the first half of the time series to the second. And, any variation at higher frequencies will come from shocks in successive periods coincidentally canceling each other out, with the likelihood of longer and longer coincidences decreases once again at an exponential rate.

In Figure 5, the fact white noise and random walks are dominated by fluctuations at polar opposite timescales manifests itself as a giant "X". The log share of trading-volume variance has a slope of $+1$ for white noise; whereas, the slope is -1 for a random walk. The center panel of Figure 5 then shows how to move between these two extremes by adjusting the persistence parameter, ρ, in an AR(1) model:

$$x(t) = \rho \cdot x(t-1) + z(t) \quad \text{where} \quad z(t) \overset{\text{iid}}{\sim} \mathrm{N}(0,1)$$

When $\rho = 0$, the process is white noise; whereas, when $\rho = 1$, it is a random walk. Thus, increasing the persistence parameter will lower the slope of the log share of trading-volume variance. And, because differences in AR(1) persistence can be converted into differences in mean-reversion speeds, which have units of 1/time, this conversion will be useful in our empirical analysis below.

Recall that we introduced the wavelet-variance estimator as a way of correcting instabilities that emerge because trading-volume variance is inherently non-stationary. So, it would

be nice to confirm that the wavelet-variance estimator is not affected by the introduction of stochastic volatility. We do this by simulating data from both an ARCH(1) model and a GARCH(1, 1) model:

$$x(t) = \sigma(t) \cdot z(t) \quad \text{with} \quad \sigma(t)^2 = \begin{cases} 0.25 + 0.90 \cdot z(t-1)^2 & \text{ARCH(1)} \\ 0.25 + 0.45 \cdot \sigma(t-1)^2 + 0.45 \cdot z(t-1)^2 & \text{GARCH(1,1)} \end{cases}$$

Without the stochastic volatility, both of these data-generating processes would just be white noise. So, we should find that they have wavelet-variance decompositions with a slope close to +1. This is exactly what we find in the right panel of Figure 5.

Confidence Intervals. We can also compute confidence intervals for the wavelet-variance estimator using standard statistical techniques (Percival, 1995). These confidence intervals will make it easier for us to evaluate the statistical significance of our main results below.

To see how, note that for any variable $x(t)$ with mean $\mathrm{E}(x) = \mu$ and variance $\mathrm{Var}(x) = \sigma^2$, the confidence interval for the sample mean, $\bar{x} \stackrel{\text{def}}{=} \frac{1}{T} \cdot \sum_{t=1}^{T} x(t)$, is:

$$\bar{x} \pm \mathrm{q}(\alpha) \cdot \hat{\sigma} / \sqrt{T-1}$$

In the equation above, the sample variance is $\hat{\sigma}^2 \stackrel{\text{def}}{=} \frac{1}{T-1} \cdot \sum_{t=1}^{T} \{x(t) - \bar{x}\}^2$, the factor of $T-1$ denotes the number of degrees of freedom when estimating the sample variance, and the function $\mathrm{q}(\cdot)$ is just the quantile function for the standard-normal distribution:

$$\Pr\left[\, |z| \leq \mathrm{q}(\alpha) \,\right] = 1 - \alpha \quad \text{where} \quad z(t) \stackrel{\text{iid}}{\sim} \mathrm{N}(0,1)$$

And, looking at Equation 5, it is clear that the wavelet-variance estimator at frequency f is the sample mean of the squared wavelet coefficients at frequency f. So, we can compute the confidence interval for WVar_f using the same approach. The sample variance is the variance of the squared wavelet coefficients at frequency f:

$$\hat{\sigma}^2 = \tfrac{1}{2^{f-1}-1} \cdot \textstyle\sum_{\ell=1}^{2^{f-1}} \{\hat{\theta}_{f,\ell} - \bar{\theta}_f\}^2$$

The number of degrees of freedom is the number of wavelets at frequency f minus one, $2^{f-1} - 1$. And, the quantile function remains unchanged.

3 Descriptive Statistics

Due to data limitations, previous research on trader horizons has been restricted to studying either a small subset of stocks or a narrow range of horizons. By contrast, you only need publicly available trading-volume data to use the wavelet-variance estimator. So, in this section, we document how the distribution of traders' investment horizons varies across all NYSE-listed stocks.

3.1 Variable Construction

We begin by describing our data and defining our key variable of interest.

Data Description. We apply the wavelet-variance estimator to minute-level trading-volume data for NYSE-listed stocks during regular trading hours, 9:30am to 4:00pm. The data come from the Trade and Quote (TAQ) database. Our sample consists of the $M = 108$ months from January 2002 to December 2010. Our sample begins in 2002 because this was the first full year with decimal pricing and because the New York Stock Exchange launched OpenBook on January 24th, 2002.

We apply the standard restrictions to this data. We remove stocks that have a price less than \$5 on the last day of the previous month as well as stocks in the bottom 30% of market capitalization on the last day of the previous month. We also only use stocks that have at least 17 days of data each month. Together, these restrictions leave us with around 1,000 stocks in each month during our sample period.

To get a sense of the timescales involved, note that for each stock there are $\frac{60 \text{ min}}{1 \text{ hr}} \times \frac{6.5 \text{ hr}}{1 \text{ day}} = 390$ minute-level trading-volume observations per day and around $\frac{21 \text{ days}}{1 \text{ month}} \times \frac{390 \text{ min}}{1 \text{ day}} = 8{,}190$ minute-level observations per month. Thus, fluctuations in volume at the 1-minute timescale roughly correspond to fluctuations at frequency $f = 13$ since $2^{13} = 8{,}192$. We account for the fact that the number of observations may not be an integer power of 2 by using the maximal-overlap discrete wavelet transform (MODWT; Percival, 1995), which is a modified version of the wavelet-transform described above.

Short-Run Tilt. When we apply the wavelet-variance estimator to a stock's trading-volume data in a given month, we get back a list of numbers corresponding to the fraction of trading-volume variance at each frequency $f \in \{2, \ldots, 13\}$. In other words, we get back an entire distribution. Distributions can be clumsy to work with empirically, so we define a new variable—namely, a stock's "short-run tilt"—as a way of summarizing the information generated by the wavelet-variance estimator in a single number.

Definition 3 (Short-Run Tilt)**.** *Stock n's short-run tilt in month m, $\texttt{tilt}_n(m)$, is the slope coefficient, $\hat{\beta}_n(m)$, from a regression of log wavelet-variance fraction on frequency:*

$$\log_2 \text{WFrac}_f(vlm_n \mid t \in m) = \hat{\alpha}_n(m) + \hat{\beta}_n(m) \cdot f + \epsilon_{n,f}(m) \tag{9}$$

We estimate the coefficient $\hat{\beta}_n(m)$ separately for each stock in our sample every month.

This definition is motivated by our analysis in Section 2.5 where we saw that white noise is dominated by high-frequency fluctuations and random walks are dominated by low-frequency fluctuations. The left panel of Figure 5 shows that the relationship between frequency and

the log wavelet-variance fraction succinctly captures this difference in timescales: a white-noise process has a slope of $+1$ while a random walk has a slope of -1. Using short-run tilt to summarize a stock's wavelet-variance decomposition is similar in spirit to using a "slope" factor to summarize the shape of the yield curve (Ang and Piazzesi, 2003).

3.2 Cross-Sectional Dispersion

Let's now look at how short-run tilt varies across NYSE-listed stock.

Portfolio Averages. To measure how the distribution of investment horizons varies across NYSE-listed stocks each month, we sort stocks into 5 value-weighted portfolios by short-run tilt. The left panel of Figure 6 reports the average short-run tilt of each of these portfolios over the $M = 108$ months in our sample period. We find that the typical stock in the highest quintile has a short-run tilt of 0.86 while the typical stock in the lowest quintile has a short-run tilt of only 0.44. We represent the 0.42-point difference between these average short-run tilts with a solid black bar.

This 0.42-point difference is well outside what we should expect to find due to measurement error. To show this, we generate 10,000 months of minute-level data for $N = 1{,}000$ stocks where each stock's demeaned trading-volume data comes from a white-noise process, $vlm(t) - \overline{vlm} \overset{\text{iid}}{\sim} \mathrm{N}(0, 1)$. Then, we sort the stocks within in each of these simulated months into quintiles based on their estimated short-run tilt and compute the realized difference in average tilts between the highest and lowest quintiles. The right panel of Figure 8 shows that on average this difference was only 0.13 points and that it was less than 0.22 points in 95% of the time. In fact, the difference never exceeded 0.30 in any of our 10,000 simulations.

Economic Magnitude. But, what is the economic meaning of 0.42? To answer this question, we use the analysis in Section 2.5. The right panel of Figure 6 depicts the persistence parameter, ρ, of an AR(1) process with the same short-run tilt as each of the 5 portfolios. If each stock's demeaned trading-volume data came from an AR(1) process, $vlm(t) = (1-\rho) \cdot \overline{vlm} + \rho \cdot vlm(t-1) + z(t)$ with $z(t) \overset{\text{iid}}{\sim} \mathrm{N}(0, \sigma^2)$, then the typical stock in the highest short-run tilt quintile would have a persistence parameter of $\rho = 0.41$ while the typical stock in the lowest quintile would have a persistence parameter of $\rho = 0.88$.

These point estimates are useful because they allow us to interpret differences in short-run tilt as differences in mean-reversion timescales. An AR(1) process with a persistence parameter of $\rho = 0.41$ has a mean-reversion timescale of $1/\{1 - 0.41\} = 1.69$ minutes while a process with a persistence parameter of $\rho = 0.88$ has a mean-reversion timescale of $1/\{1 - 0.88\} = 8.33$ minutes. In other words, increasing a stock's short-run tilt from 0.44 to

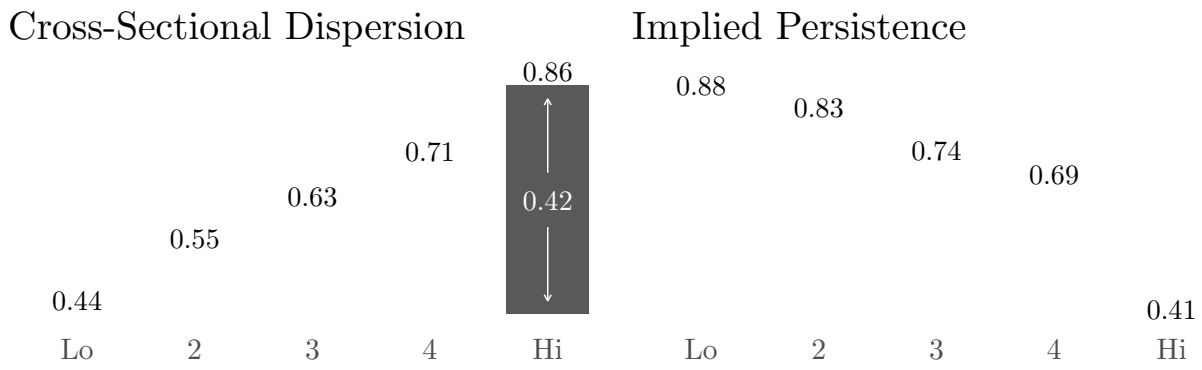

Figure 6: *Estimated values for value-weighted monthly portfolios sorted by short-run tilt.* **Cross-Sectional Dispersion.** *Numbers and height depict the average short-run tilt,* $\overline{\texttt{tilt}}$*, of the stocks in each portfolio over the* $M = 108$ *months in our sample. Solid bar depicts the difference between averages of the "Lo" and "Hi" portfolios. Reads: "A stock in the highest quintile typically has a short-run tilt of* 0.86*, which is* 0.42 *higher than the typical short-run tilt of a stock in the lowest quintile."* **Implied Persistence.** *Numbers and height depict the persistence parameter,* ρ*, of an* AR(1) *process with the same tilt as each portfolio. Reads: "If each stock's trading-volume data came from an* AR(1) *process, then the typical stock in the highest short-run tilt quintile would have a persistence parameter of* $\rho = 0.41$ *while the typical stock in the lowest quintile would have a persistence parameter of* $\rho = 0.88$*. And, these persistence parameters imply mean-reversion timescales that differ by a factor of* 5*:* $1/\{1 - 0.41\} = 1.69$ *minutes for stocks in the highest quintile and* $1/\{1 - 0.88\} = 8.33$ *minutes for stocks in the lowest quintile."*

0.86 corresponds to trading-volume data that mean reverts 5-times faster. So, stocks in the lowest quintile have traders with investment horizons that are 5-times longer than those of stocks in the highest quintile.

Control Variables. Of course, it could be that this cross-sectional dispersion in short-run tilt is the result of cross-sectional differences in some other variable. To investigate this possibility, we look at how firm characteristics, liquidity, and news announcements vary across portfolios sorted by short-run tilt. The top row of Figure 7 depicts results for market cap, volume, and implied volatility. We find that stocks with a larger share of their trading volume-variance at higher frequencies tend to be larger stocks with less volatile daily returns. But, these stocks are not traded significantly more often.

The middle row of Figure 7 shows that the stocks in the highest short-run tilt quintile also tend to be slightly less liquid. We proxy for this illiquidity in two different ways. First, we use the Amihud (2002) measure, which defines a stock's illiquidity in a given month as the daily average absolute return per dollar traded. Second, we compute the daily average bid-ask spread as a percent of the closing price for each stock during a month.

Finally, the remaining panels in Figure 7 show that the stocks in the highest short-run tilt quintile are no more likely to realize news about long-term fundamentals than stocks in

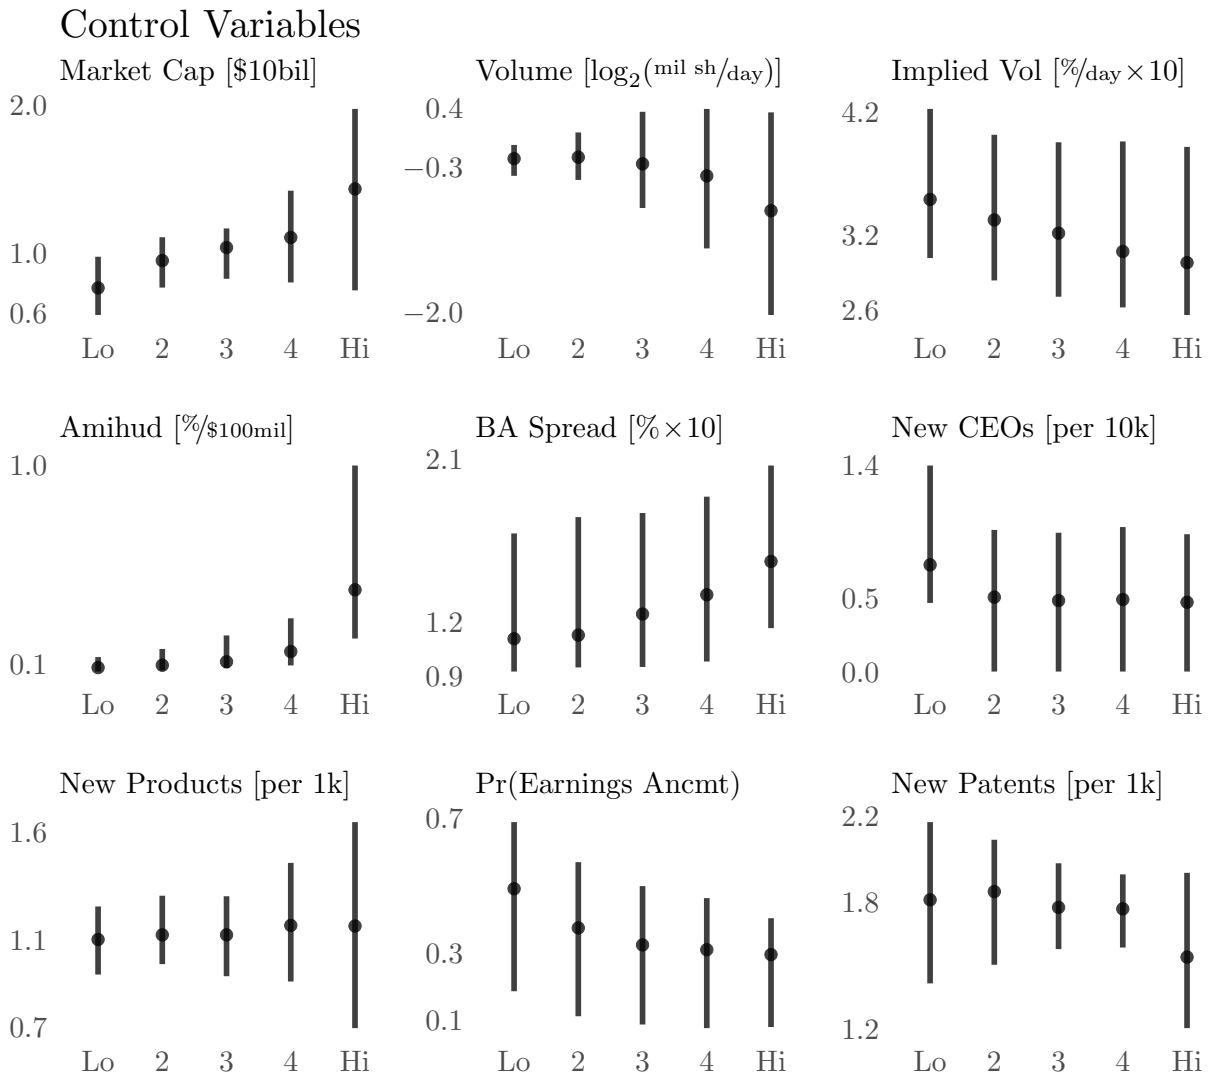

Figure 7: *Summary statistics describing the value of control variables for portfolios sorted by short-run tilt each month. Dots represent portfolio averages over the $M = 108$ months in our sample. Vertical lines represent inter-quartile ranges.* **Market Cap.** *Market capitalization at end of previous month; +1 represents an additional $10bil in market cap.* **Volume.** *Average daily trading volume; +1 represents double the average trading volume.* **Implied Vol.** *Average daily implied volatility; +1 represents a 0.1% increase in daily volatility.* **Amihud.** *Average absolute return per dollar traded each day; +1 represents an additional 1% per $100mil traded.* **BA Spread.** *Average bid-ask spread at market close as percent of closing price; +1 represents a 0.1% increase in bid-ask spread.* **New CEOs.** *Number of CEO changes in month per 10,000 firms; +1 represents a 0.01% increase in probability.* **New Products.** *Number of product announcements in month per 1,000 firms; +1 represents a 0.1% increase in probability.* **Pr(Earnings Ann.)** *Probability that stock has earnings announcement in month.* **New Patents.** *Number of patent grants in month per 1,000 firms; +1 represents a 0.1% increase in probability.*

the lowest quintile. We look at four different kinds of news announcements. First, we use the ExecuComp database to identify the date of CEO turnovers (Pan et al., 2015). Second, we get the date of new-product announcements from the Key Development Database of Capital IQ (Liu and Xuan, 2017). Third, we get earnings-announcement dates from the I/B/E/S database (Della Vigna and Pollet, 2009). And fourth, we use the patent-grant data for firms in CRSP from Kogan et al. (2017). We find no statistically significant differences in any of these announcement-related variables across our short-run tilt portfolios. And, it would be surprising if we did. As emphasized in the introduction, there are not enough announcements about long-term fundamentals to explain minute-to-minute differences in volume.

Sorting on Residuals. But, none of these control variables explain the cross-sectional dispersion in short-run tilt that we observe in the data. To show that this is the case, we regress each stock's short-run tilt in a given month on each of the above control variables and its squared value:

$$\texttt{tilt}_n(m) = \hat{\alpha}(m) + \hat{\beta}(m) \cdot \texttt{x}_n(m) + \hat{\gamma}(m) \cdot \texttt{x}_n(m)^2 + \texttt{e}_n(m) \tag{10}$$

The residuals from each of these $M = 108$ regressions, $\texttt{e}_n(m)$, then represent the portion of a stock's short-run tilt that is unrelated to cross-sectional differences in a particular control variable. We include the squared terms to account for any non-linear variations across portfolios. The left panel of Figure 8 shows the difference between the value-weighted average short-run tilt of portfolios sorted by these regression residuals. The amount of cross-sectional dispersion hardly changes when sorting portfolios on residual rather than actual short-run tilt. There is still a statistically significant difference between the short-run tilt of the highest and lowest quintiles after controlling for firm characteristics, liquidity, and news announcements.

3.3 Measure Verification

We also show that short-run tilt is consistent with the measure of short-run trading in Brogaard et al. (2014). This earlier estimate of short-run trading activity is more direct, but it can only be computed for a subset of 120 stocks during 2008, 2009, and 2010.

Variable Construction. The NASDAQ uses internal data to categorize each trader registered with the exchange as either a high-frequency trader (HFT) or a non-HFT. For example, the NASDAQ tracks how often a trader's net intraday position crosses zero, the length of a trader's average order duration, and his order-to-trade ratio. And, for a subset of 120 stocks (60 of which are listed on the NYSE) the NASDAQ released data on HFT order flow for 2008, 2009, and 2010. For stocks in this sample, the exchange released data for each stock

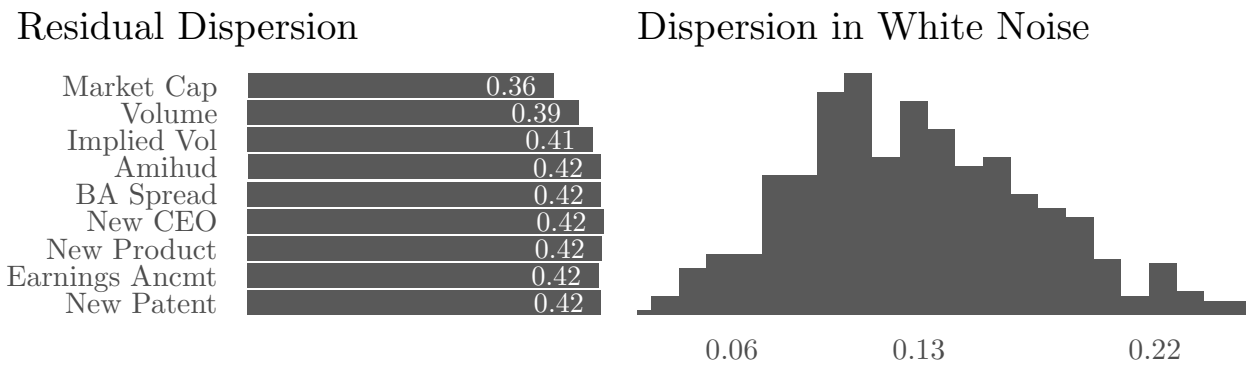

Figure 8: Residual Dispersion. *Bars represent the average difference between the short-run tilt in the highest quintile and the short-run tilt in the lowest quintile when sorting stocks into value-weighted portfolios based on their residual short-run tilt as defined in Equation 10. Reads: "There is still a statistically significant difference between the short-run tilt of the highest and lowest quintiles after controlling for firm characteristics, liquidity, and news announcements."* **Dispersion in White Noise.** *Histogram showing the average cross-sectional dispersion in white noise. This distribution comes from 10,000 simulations where we sort* $N = 1{,}000$ *stocks into quintiles based on their short-run tilt in a given month and each stock's demeaned trading volume is generated by a white-noise process,* $vlm(t) - \overline{vlm} \overset{\text{iid}}{\sim} \mathrm{N}(0,1)$. *The* x*-axis represents the realized difference in short-run tilt between the highest and lowest quintiles. The numbers on the* x*-axis denote the mean and* 95% *confidence interval from the 10,000 simulations. The bar heights represent the probability that the high-minus-low difference is equal to* x *in a randomly selected simulation. We do not label the* y*-axis because the values depend on the bin width, and there is no preferred bin width. Reads: "In* 95% *of simulations using white-noise data, the average difference between the short-run tilt in the highest and lowest quintiles is less than* 0.22*."*

on the volume coming from trades that involve two HFTs, $\texttt{both}_n(m)$, the volume coming from trades that involve an HFT and a non-HFT, $\texttt{either}_n(m)$, and the volume coming from trades that involve two non-HFTs, $\texttt{neither}_n(m)$.

Brogaard et al. (2014) used this data to estimate the fraction of each stock's trading volume in a given month that was due to HFTs. To start with, they estimate this fraction as the share of a stock's volume, $\texttt{vlm}_n(m)$, that came from HFTs trading with one another:

$$\texttt{onlyHft}_n(m) \overset{\text{def}}{=} \frac{\texttt{both}_n(m)}{\texttt{vlm}_n(m)} \tag{11}$$

But, they also estimate this fraction as the share of a stock's volume that involves HFTs in any capacity:

$$\texttt{anyHft}_n(m) \overset{\text{def}}{=} \frac{\texttt{both}_n(m) + \texttt{either}_n(m)}{\texttt{vlm}_n(m)} \tag{12}$$

The top panel of Table 1 gives summary statistics for both of these measures.

Regression Results. The bottom panel of Table 1 then reports the estimated coefficients

from a regression of short-run tilt on these two measures of HFT trading activity:

$$\widetilde{\texttt{tilt}}_n(m) = \hat{\beta} \cdot \widetilde{\texttt{onlyHft}}_n(m) + \texttt{e}_n(m) \tag{13a}$$

$$\widetilde{\texttt{tilt}}_n(m) = \hat{\beta} \cdot \widetilde{\texttt{anyHft}}_n(m) + \texttt{e}_n(m) \tag{13b}$$

We normalize all variables in the regression to have zero mean and unit standard deviation, $\tilde{x} \stackrel{\text{def}}{=} {}^{\{x-\mu\}}/_{\sigma}$, to make the coefficients easier to interpret. The first column shows that a 1-standard-deviation increase in the fraction of a stock's trading volume in month m that came from HFTs trading with one another is associated with a 0.22-standard-deviation increase in a stock's short-run tilt. The second column shows that this result remains when we include lagged values of each of these variables. The third and fourth columns show that this finding is robust to using the alternative measure of HFT share.

Key Differences. While it is reassuring to know that short-run tilt is positively related to Brogaard et al. (2014)'s estimate of HFT share, there are three reasons why we would not expect to find a perfect correlation. First, the NASDAQ classifies HFTs at the firm level. So, HFT desks in large institutions, such as Goldman Sachs and Morgan Stanley, will be excluded. Small HFTs that route their orders through these large institutions will also be ignored. Second, the NASDAQ's classification does not vary over time. So, if some low-frequency firms were to gradually speed up over time, the NASDAQ's measure would understate the HFT share later in the sample. Third, HFTs can trade much faster than once a minute. So, HFT share and short-run tilt measure activity at slightly different timescales.

4 Horizon Spillover

Now that we have described our new statistical tool and provided descriptive statistics, we can present our two main findings.

4.1 Abnormal Returns

First, we show that stocks dominated by short-run fluctuations in volume have abnormal returns that are 1% per month higher than otherwise similar stocks with less short-run tilt.

Result #1. This point estimate comes from regressing the monthly excess returns of 5 portfolios sorted by short-run tilt on the Fama and French (1993) factors:

$$\texttt{rx}_i(m) = \hat{\alpha} + \hat{\beta}_{\text{Mkt}} \cdot \texttt{rx}_{\text{Mkt}}(m) + \hat{\beta}_{\text{SmB}} \cdot \texttt{rx}_{\text{SmB}}(m) + \hat{\beta}_{\text{HmL}} \cdot \texttt{rx}_{\text{HmL}}(m) + \texttt{e}_i(m) \tag{14}$$

If the excess returns earned by the quintile of stocks with the largest short-run tilt were explained by variation in exposures to the Fama and French (1993) risk factors, then the $\hat{\alpha}$

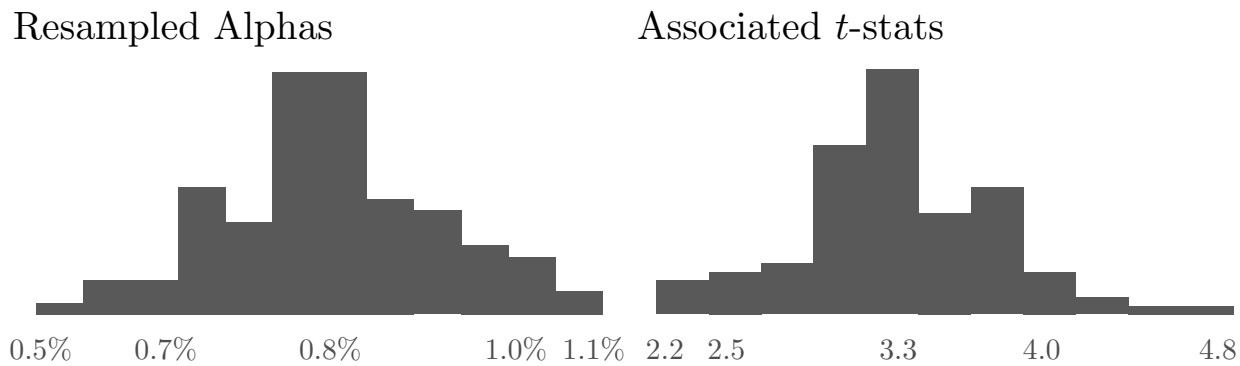

Figure 9: *Results of resampling procedure to assess effects of measurement error. Single iteration of the procedure: a) For each stock n in month m, we compute the* 95% *confidence interval of the wavelet-variance estimator at frequency f. b) Then, we draw a new estimate for the wavelet-variance fraction at frequency f from a uniform distribution over this* 95% *confidence interval. c) We repeat this process at every frequency and use the results to estimate new value for stock n's short-run tilt in month m. d) We re-estimate the abnormal returns to sorting by short-run tilt reported in the top panel of Table 2 using the resampled values for every stock-month. We iterate on this procedure 1,000 times.* **Resampled Alphas.** *The x-axis represents monthly alphas. Numbers on the x-axis denote the min,* 5%*tile, mean,* 95%*tile, and max over the 1,000 iterations. Bar heights represent probability that the resampled "Hi-Lo" $\hat{\alpha}$ equals x. There is no y-axis label because values depend on the histogram bin width, and there is no preferred bin width.* **Associated t-stats.** *The x-axis represents the associated t-stats. Numbers on the x-axis denote the min,* 5%*tile, mean,* 95%*tile, and max from the 1,000 iterations. Bar heights represent probability that the resampled "Hi-Lo" $\hat{\alpha}$ has a t-stat equal to x. There is no y-axis label because values depend on the histogram bin width, and there is no preferred bin width. Reads: "The abnormal returns to sorting on a noisy version of short-run tilt are lower than in top panel of Table 2, but even with the additional noise these abnormal returns still have an average t-stat of* 3.3*."*

in Equation 14 should be zero for each portfolio. This is not what we find in the top panel of Table 2. The quintile of stocks with the largest short-run tilt has an $\hat{\alpha} = 0.32\%$ per month while the quintile of stocks with the smallest short-run tilt has an $\hat{\alpha} = -0.65\%$ per month. And, the corresponding high-minus-low portfolio yields an $\hat{\alpha} = 0.97\%$ per month, which is both statistically and economically significant. Thus, differences in short-run tilt are priced in monthly returns. The relative share of trading activity at the 1-minute timescale matters for investors at the 1-month horizon.

Robustness Checks. And, Table 3 shows that these abnormal returns are not explained by firm characteristics, liquidity, momentum, or news announcements. This table contains two kinds of tests depending on the control variable. If the variable is an excess return, then we just include it as an additional factor in the regression from Equation 14. Otherwise, we estimate the Fama and French (1993) abnormal returns of portfolios sorted by the residual tilt relative to this variable as in Hou and Loh (2014).

The first 3 rows of Table 3 show the abnormal returns associated with short-run tilt are not due to size, volume, or implied volatility. The next 5 rows look at the relationship between these abnormal returns and various liquidity measures. In addition to the Amihud (2002) measure and the bid-ask spread, we also consider the Pastor and Stambaugh (2003) liquidity factor, exposure to daily liquidity shocks (Acharya and Pedersen, 2005), and separate exposures to the permanent and transient components of these shocks (Sadka, 2006). We find no evidence that the abnormal returns associated with short-run tilt are related to any of these liquidity measures.

If a stock performs very well, then more traders will have to add or remove it from their portfolio. So, perhaps short-run tilt is driven by a stock's past performance? The next 2 rows in Table 3 suggest that this is not the case. They show that the abnormal returns associated with short-run tilt persist after including the Carhart (1997) momentum factor and after controlling for a stock's excess return in the previous month. And, the final 4 rows show that the abnormal returns are not due to cross-portfolio differences in the amount of news about long-term fundamentals.

Measurement Error. Short-run tilt is an estimated variable representing a true value plus some measurement error. And, when we compute the abnormal returns to portfolios sorted by short-run tilt, this measurement error will distort our abnormal-return estimates. But, using the confidence intervals calculated in Section 2.5, we estimate that these distortions are unlikely to have a significant effect on the standard errors reported in Table 2.

Our approach is similar to bootstrapping. For each stock n in month m, we can compute the 95% confidence interval of the wavelet-variance estimator at frequency f. We can use this information to draw a new estimate for the wavelet-variance fraction at frequency f from a uniform distribution with upper and lower bounds defined by the 95% confidence interval. If we repeat this process for every frequency, then we can use the results to compute a new estimate for stock n's short-run tilt in month m. And, if we do this for every stock in every month, then we can re-estimate the abnormal returns to sorting by short-run tilt reported in the top panel of Table 2.

Figure 9 shows the distribution of the estimated $\hat{\alpha}$ (left panel) and the associated t-stats (right panel) for the "Hi-Lo" portfolio in 1,000 iterations of this resampling procedure. Adding additional measurement error to our best estimates of short-run tilt values should lower our cross-sectional $\hat{\alpha}$. And, in the left panel of Figure 9, this is exactly what we see. The abnormal returns to sorting on a noisy version of short-run tilt drop to 0.8% per month. However, even with the additional noise, the right panel shows that these abnormal returns still have an average t-stat of 3.3. And, in none of the 1,000 resamplings did the t-stat on the "Hi-Lo" portfolio drop below 2.2.

4.2 Rapid Churn

Second, we show that a stock's short-run tilt can change rapidly from month to month, supporting the idea that these abnormal returns are not due to some persistent firm characteristic that's adding both short-run fluctuations and long-term risk. We view this rapid churn as a new stylized fact that's worthy of further investigation.

Result #2. The bottom panel of Table 2 presents this result. The top left entry in this matrix says that, if a stock is in the quintile with the largest short-run tilt in month m, there is only a 32% chance that the stock is still in the quintile with the largest short-run tilt in month $\{m+1\}$. And, the top right entry says that roughly 1 out of every 10 stocks in the quintile with the largest short-run tilt in month m will belong to the quintile with the smallest short-run tilt in month $\{m+1\}$. Scanning down the rows in this matrix, we can also see that the effect is roughly symmetric, with stocks in the smallest quintile in month m moving to the largest quintile in month $\{m+1\}$ at roughly the same rate.

This rapid churn implies that control variables computed over the course of several months are going to do a poor job of explaining abnormal returns earned associated with short-run tilt. So, this observation is entirely consistent with the results in Table 3, which finds almost no link between the abnormal returns to short-run tilt and firm characteristics, liquidity measures, past performance, or news announcements. This finding is also related to Lo and Wang (2000)'s finding that the level of trading volume is not very persistent.

Trading Strategy. The observation also implies that it should only be possible to trade on the information in short-run tilt at horizons shorter than a month. This is exactly what we find in Table 4. We study the abnormal returns in the first 1, 2, 3, and 4 weeks of month m of a trading strategy that is long the quintile of stocks with the largest short-run tilt in month $\{m-1\}$ and short the quintile of stocks with the smallest short-run tilt in month $\{m-1\}$. When the holding period is only 1 week, the strategy generates an abnormal return of 0.93% per month relative to the Fama and French (1993) 3-factor model. But, when we extend the holding period to 2 weeks, this point estimate drops to 0.29% per month, too small to survive transaction costs. With a 3-week holding period, it drops even further to 0.14% per month. And, when we hold the portfolio for a full month, the abnormal returns become statistically indistinguishable from zero.

4.3 Mechanism

Although it helps make sense of our regression results in Table 3 and trading-strategy returns in Table 4, this rapid churn also presents a new puzzle of its own. Namely, if short-run tilt

changes rapidly from month to month, how do long-term investors know which stocks have large tilts? This subsection presents evidence for one plausible mechanism: execution risk. Although an investor trading once a day, for example, cannot interpret minute-to-minute fluctuations in volume, he can track his end-of-day order-execution statistics. We find that half of the abnormal returns to short-run tilt are associated to differences in execution risk.

Variable Construction. To compute our measure of execution risk for long-term investors, we use data from Anand et al. (2012) on institutional equity transactions compiled by Ancerno Ltd. The dataset contains information on individual tickets sent by institutions to their broker. Ancerno is a consulting firm that compiles this information and then reports end-of-day summary statistics to their institutional clients.

Let $k = 1, \ldots, K_n$ count the tickets sent to a broker by any institution with the goal of trading stock n. Let $q_n(k)$ denote the last quoted price of stock n prior to the kth ticket, and let $p_n(k)$ denote the average price at which the kth ticket was executed. We define the execution shortfall associated with the kth ticket for stock n as the relative difference between the average execution price and the initial price where $buy_n(k)$ and $sell_n(k)$ are indicator variables for whether the kth ticket was a buy or sell order:

$$shortfall_n(k) \stackrel{\text{def}}{=} \left\{ \frac{p_n(k) - q_n(k)}{q_n(k)} \right\} \times \{buy_n(k) - sell_n(k)\}$$

We then define the execution risk associated with stock n in month m as the standard deviation of this shortfall for tickets submitted in month m:

$$\texttt{execRisk}_n(m) \stackrel{\text{def}}{=} \text{Sd}(shortfall_n | k \in m) \tag{15}$$

A higher standard deviation means that an institutional investor who wants to buy a share of stock n in month m will be less certain about the difference between the price he is currently being quoted and his eventual execution price.

Abnormal Returns. To see how much of the abnormal returns to short-run tilt operate through execution risk, we follow our earlier approach and estimate the Fama and French (1993) abnormal returns of portfolios sorted by the residual tilt relative to execution risk as defined in Equation 10. Table 5 reports that, after accounting for cross-sectional variation in execution risk, the abnormal returns associated with short-run tilt drop to 0.47% per month. And, this reduction in abnormal returns is consistent across portfolios.

Execution Risk vs. Liquidity Risk. Execution risk is a type of liquidity risk. However, we already tested whether liquidity risk as defined in Pastor and Stambaugh (2003) and Acharya and Pedersen (2005) explained the abnormal returns associated with short-run tilt in Table 3, and we found that it did not. But, these two earlier measures differ from execution risk in

their level of granularity. In Pastor and Stambaugh (2003), stocks with lots of liquidity risk in a month have lots of day-to-day return reversals. And, in Acharya and Pedersen (2005), a stock's liquidity risk is measured as its covariance over the past 18 months with the Amihud (2002) illiquidity measure, which is a daily average. So, both of these variables are daily measures of liquidity risk within a month.

But, stocks with a large short-run tilt are the stocks where a daily measure of liquidity risk is the least appropriate. These are the stocks where traders have the shortest horizons. For these stocks, what a monthly investor really cares about is a stock's liquidity at the exact time he decides to place an order. This is what we are measuring with the execution risk defined using Ancerno's ticket-level dataset. And, we know that longer-term investors see this data at the end of each day because this is exactly the data that Ancerno is in the business of providing. The fact that the abnormal returns to short-run tilt are related to execution risk but not daily measures of liquidity risk suggests that the way you measure liquidity risk in empirical applications should depend on the horizon of the traders involved.

5 Conclusion

This paper studies how short-run trading activity affects long-term returns by decomposing each stock's trading-volume variance into frequency-specific components using the wavelet-variance estimator (Percival, 1995). After introducing this new statistical tool, we present two main results. First, we show that stocks dominated by short-run fluctuations in volume have abnormal returns that are 1% per month higher than otherwise similar stocks with less of a short-run tilt. Second, we show that a stock's short-run tilt can change rapidly from month to month, suggesting these abnormal returns aren't due to some persistent firm characteristic that's adding both short-run fluctuations and long-term risk.

This paper also makes a more general methodological contribution. The ability to measure the entire distribution of investment horizons for the entire cross-section of stocks using only publicly available data has a wide range of potential applications. For instance, one of the key constraints that arbitrageurs face in the limits-to-arbitrage models is a short investment horizon (De Long et al., 1990). There is also literature looking at the optimal response of firm managers to shareholders with biased short-term beliefs (Baker and Wurgler, 2013). And, this rational-manager vs biased-shareholder paradigm fits into an even larger literature looking at feedback effects between financial markets and the real economy (Chen et al., 2007) where the decisions of myopic investors also play a key role.

References

Acharya, V. and L. Pedersen (2005). Asset pricing with liquidity risk. *Journal of Financial Economics 77*(2), 375–410.

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

Anand, A., P. Irvine, A. Puckett, and K. Venkateraman (2012). Performance of institutional trading desk: An analysis of persistence in trading costs. *Review of Financial Studies 25*(2), 557–598.

Andersen, T. (1996). Return volatility and trading volume: An information flow interpretation of stochastic volatility. *Journal of Finance 51*(1), 169–204.

Ang, A. and Piazzesi (2003). A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. *Journal of Monetary Economics 50*(4), 745–787.

Baker, M. and J. Wurgler (2013). Behavioral corporate finance: An updated survey. *Handbook of the Economics of Finance 2*, 357–424.

Bandi, F., B. Perron, A. Tamoni, and C. Tebaldi (2015). The scale of predictability. Working Paper.

Beber, A., J. Driessen, and P. Tuijp (2012). Pricing liquidity risk with heterogeneous investment horizons. Working Paper.

van Binsbergen, J., M. Brandt, and R. Koijen (2012). On the timing and pricing of dividends. *American Economic Review 102*(4), 1596–1618.

Brogaard, J., T. Hendershott, and R. Riordan (2014). High-frequency trading and price discovery. *Review of Financial Studies 27*(8), 2267–2306.

Campbell, J., S. Grossman, and J. Wang (1993). Trading volume and serial correlation in stock returns. *Quarterly Journal of Economics 108*(4), 905–939.

Carhart, M. (1997). On persistence in mutual fund performance. *Journal of Finance 52*(1), 57–87.

Cella, C., A. Ellul, and M. Giannetti (2013). Investors' horizons and the amplification of market shocks. *Review of Financial Studies 26*(7), 1607–1648.

Chen, Q., I. Goldstein, and W. Jiang (2007). Price informativeness and investment sensitivity to stock price. *Review of Financial Studies 20*(3), 619–650.

Crouzet, N., I. Dew-Becker, and C. Nathanson (2017). Multi-frequency trade. Working Paper.

De Long, B., A. Shleifer, S. Lawrence, and R. Waldmann (1990). Noise trader risk in financial markets. *Journal of Political Economy 98*(4), 703–738.

Della Vigna, S. and J. Pollet (2009). Investor inattention and friday earnings. *Journal of Finance 64*(2), 709–749.

Dew-Becker, I. and S. Giglio (2016). Asset pricing in the frequency domain: Theory and empirics. *Review of Financial Studies Forthcoming.*

Easley, D., M. López de Prado, and M. O'Hara (2012). The volume clock: Insights into the high-frequency paradigm. *Journal of Portfolio Management 39*(1), 19–29.

Epps, T. and M. Epps (1976). The stochastic dependence of security price changes and transaction volumes: Implications for the mixture-of-distributions hypothesis. *Econometrica 44*(2), 305–321.

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

Gallant, R., P. Rossi, and G. Tauchen (1992). Stock prices and volume. *Review of Financial Studies 5*(2), 199–242.

Gençay, R., F. Selçuk, and B. Whitcher (2001). *An Introduction to Wavelets and Other Filtering Methods in Finance and Economics* (1 ed.). Academic Press.

Harris, L. (1987). Transaction data test of the mixture of distributions hypothesis. *Journal of Financial and Quantitative Analysis 22*(1), 127–141.

Hasbrouck, J. (2015). High-frequency quoting: Short-term volatility in bids and offers. SSRN Working Paper.

Hou, K. and R. Loh (2014). Have we solved the idiosyncratic volatility puzzle? Working Paper.

Jagannathan, R. and Y. Wang (2007). Lazy investors, discretionary consumption, and the cross-section of stock returns. *Journal of Finance 62*(4), 1623–1661.

Kamara, A., R. Korajczyk, X. Lou, and R. Sadka (2015). Horizon pricing. *Journal of Financial and Quantitative Analysis (Forthcoming).*

Kirilenko, A., A. Kyle, M. Samadi, and T. Tuzun (2011). The flash crash: The impact of high frequency trading on an electronic market. Working paper.

Kogan, L., D. Papnikolaou, A. Seru, and N. Stoffman (2017). Technological innovation, resource allocation, and growth. *Quarterly Journal of Economics 132*(2), 665–712.

Kyle, A. (1985). Continuous auctions and insider trading. *Econometrica 53*(6), 1315–1335.

Lan, C., F. Moneta, and R. Wermers (2015). Mutual fund investment horizon and performance. Working Paper.

Liu, P. and Y. Xuan (2017). The contract year phenomenon in the corner office: An analysis of firm behavior during ceo contract renewals. Working Paper.

Lo, A. and J. Wang (2000). Trading volume: Definitions, data analysis, and implications of portfolio theory. *Review of Financial Studies 13*(2), 257–300.

MacKay, D. (2003). *Information Theory, Inference, and Learning* (1 ed.). Cambridge University Press.

Mallat, S. (1999). *A Wavelet Tour of Signal Processing* (1 ed.). Academic Press.

Pan, Y., T. Wang, and M. Weisback (2015). Learning about ceo ability and stock return volatility. *Review of Financial Studies 28*(6), 1623–1666.

Pastor, L. and R. Stambaugh (2003). Liquidity risk and expected stock returns. *Journal of Political Economy 111*(3), 642–685.

Percival, D. (1995). On estimation of the wavelet variance. *Biometrika 82*(3), 619–631.

Percival, D. (2016). A wavelet perspective on the allan variance. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control 63*(4), 538–554.

Sadka, R. (2006). Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk. *Journal of Financial Economics 80*(2), 309–349.

Serroukh, A., A. Walden, and D. Percival (2000). Statistical properties and uses of the wavelet variance estimator for the scale analysis of time series. *Journal of the American Statistical Association 95*(449), 184–196.

Vayanos, D. (2001). Strategic trading in a dynamic noisy market. *Journal of Finance 56*(1), 131–171.

Measure Verification

Summary statistics

		μ	σ	min	q_{25}	q_{50}	q_{75}	max
$\texttt{onlyHft}_n(m)$		7.73	6.80	0.00	2.32	5.24	12.11	38.82
$\texttt{anyHft}_n(m)$		49.00	18.98	0.22	33.97	48.86	65.50	87.77

Regression analysis

	$\widetilde{\texttt{tilt}}_n(m)$				
$\widetilde{\texttt{onlyHft}}_n(m)$	0.22 (0.02)	0.14 (0.06)			
$\widetilde{\texttt{onlyHft}}_n(m-1)$		0.04 (0.06)			
$\widetilde{\texttt{anyHft}}_n(m)$			0.24 (0.03)	0.00 (0.09	
$\widetilde{\texttt{anyHft}}_n(m)$				0.22 (0.09)	
$\widetilde{\texttt{tilt}}_n(m-1)$		0.24 (0.03)		0.25 (0.03)	0.29 (0.03)

Table 1: Summary statistics. *Summary statistics for NASDAQ measures of short-run trading activity computed for the available subset of* 60 *NYSE-listed stocks during 2008 and 2009.* $\texttt{onlyHft}_n(m)$ *denotes share of stock* n*'s volume in month* m *that only involves high-frequency traders (HFTs) as defined in Equation 11.* $\texttt{anyHft}_n(m)$ *denotes share of stock* n*'s volume in month* m *that involves HFTs in any capacity as defined in Equation 12. Histograms show the distribution across stock-months. Reads: "On average, for the* 60 *stocks where NASDAQ classifies trader horizons, only* 7.73% *of the trades each month only involve HFTs."* **Regression analysis.** *Coefficients from regressions of each stock's normalized short-run tilt in month* m *on normalized versions of the NASDAQ short-run trading-activity measures. A* $\widetilde{\text{tilde}}$ *denotes a normalized variable with* $\tilde{x} = (x - \mu)/\sigma$. $\widetilde{\texttt{tilt}}_n(m)$ *is the normalized value of stock* n*'s short-run tilt in month* m *as defined in Equation 9. Numbers in parentheses are standard errors. Reads: "A* 1*sd increase in the fraction of trading volume coming only from high-frequency traders is associated with a* 0.22*sd increase in a stock's short-run tilt."*

Two Main Results

Abnormal Returns

	$\hat{\alpha}$	$\hat{\beta}_{\text{Mkt}}$	$\hat{\beta}_{\text{SmB}}$	$\hat{\beta}_{\text{HmL}}$
Hi	0.32 (0.12)	0.89 (0.03)	−0.23 (0.05)	0.24 (0.05)
4	0.26 (0.13)	0.95 (0.93)	−0.23 (0.05)	0.10 (0.05)
3	0.33 (0.10)	1.01 (0.02)	−0.13 (0.04)	0.10 (0.04)
2	0.09 (0.14)	1.08 (0.03)	−0.05 (0.06)	0.15 (0.05)
Lo	−0.65 (0.18)	1.03 (0.04)	0.06 (0.08)	0.13 (0.07)
Hi − Lo	0.97 (0.24)	−0.14 (0.06)	−0.29 (0.10)	0.11 (0.09)

Transition probabilities

		Month $m+1$				
		Hi	4	3	2	Lo
Month m	Hi	0.32	0.22	0.18	0.14	0.11
	4	0.22	0.22	0.20	0.18	0.16
	3	0.18	0.20	0.20	0.20	0.20
	2	0.14	0.18	0.20	0.22	0.23
	Lo	0.11	0.16	0.20	0.23	0.27

Table 2: Abnormal returns. *Abnormal returns relative to the Fama and French (1993) factors of 5 value-weighted portfolios sorted by short-run tilt as shown in Equation 14. Sample covers the $M = 108$ months from January 2002 to December 2010. Numbers in parentheses are standard errors. Reads: "The quintile of stocks with the largest short-run tilt has abnormal returns that are 0.97% per month higher than the quintile of stocks with the smallest short-run tilt."* **Transition probabilities.** *Rows correspond to quintiles sorted by short-run tilt in month m. Columns correspond to quintiles sorted by short-run tilt in month $\{m+1\}$. Entry (i, j) in this matrix represents the probability that a stock in quintile i right now belongs to quintile j next month. Roughly 2% of stocks exit the sample due to market-cap or price restrictions each month, so the rows do not add up to 1. Reads: "If a stock belongs to the quintile with the largest short-run tilt right now, then there is still only a 32% chance that it belongs to the quintile with the largest short-run tilt next month."*

Monthly Alphas with Control Variables

	Hi	4	3	2	Lo	Hi − Lo
Market Cap	0.42 (0.14)	0.15 (0.13)	0.15 (0.11)	0.24 (0.13)	−0.65 (0.18)	1.06 (0.26)
Volume	0.35 (0.11)	0.16 (0.10)	0.31 (0.13)	0.21 (0.13)	−0.99 (0.23)	1.34 (0.28)
Implied Volatility	0.34 (0.13)	0.33 (0.13)	0.29 (0.11)	0.16 (0.14)	−0.65 (0.18)	0.99 (0.24)
Amihud	0.30 (0.12)	0.15 (0.13)	0.17 (0.09)	0.03 (0.14)	−0.68 (0.18)	0.98 (0.25)
BA Spread	0.33 (0.12)	0.20 (0.13)	0.29 (0.10)	0.12 (0.14)	−0.74 (0.19)	1.08 (0.24)
Pastor-Stambaugh	0.28 (0.12)	0.22 (0.14)	0.25 (0.10)	−0.03 (0.14)	−0.77 (0.19)	1.05 (0.25)
Acharya-Pedersen	0.29 (0.12)	0.13 (0.14)	0.25 (0.10)	−0.09 (0.13)	−0.55 (0.18)	0.85 (0.124)
Sadka	0.28 (0.12)	0.29 (0.13)	0.30 (0.11)	0.12 (0.13)	−0.60 (0.18)	0.88 (0.24)
Momentum Factor	0.32 (0.12)	0.25 (0.13)	0.30 (0.10)	0.08 (0.14)	−0.64 (0.18)	0.96 (0.24)
Lagged Returns	0.33 (0.12)	0.26 (0.13)	0.28 (0.11)	0.14 (0.14)	−0.62 (0.19)	0.95 (0.25)
New CEOs	0.32 (0.12)	0.27 (0.13)	0.32 (0.10)	0.09 (0.14)	−0.65 (0.18)	0.97 (0.24)
New Products	0.33 (0.12)	0.17 (0.12)	0.30 (0.11)	0.14 (0.13)	−0.65 (0.19)	0.99 (0.26)
Earnings Ancmt	0.37 (0.12)	0.26 (0.14)	0.33 (0.10)	0.16 (0.14)	−0.65 (0.17)	1.02 (0.24)
New Patents	0.27 (0.12)	0.25 (0.12)	0.25 (0.11)	0.15 (0.14)	−0.65 (0.18)	0.93 (0.24)

Table 3: *Monthly alphas computed during period from January 2002 to December 2010. Numbers in parentheses are standard errors. Rows labeled* **Pastor-Stambaugh** *and* **Momentum Factor** *report alphas relative to Fama and French (1993) plus either the Pastor and Stambaugh (2003) liquidity factor or the Carhart (1997) momentum factor for value-weighted portfolios sorted by short-run tilt. Remaining rows report alphas relative to just Fama and French (1993) for value-weighted portfolios sorted by residual short-run tilt as in Equation 10.* **Market Cap.** *Market cap at end of previous month.* **Volume.** *Log of average daily trading volume.* **Implied Vol.** *Average daily implied volatility.* **Amihud.** *Average absolute return per dollar traded each day (Amihud, 2002).* **BA Spread.** *Average bid-ask spread at market close as percent of closing price.* **Acharya-Pedersen.** *Exposure to variation in Amihud illiquidity during previous 18 months (Acharya and Pedersen, 2005).* **Sadka.** *Exposure to permanent and transient liquidity factors from Sadka (2006) during previous 18 months.* **Lagged Returns.** *Return in previous month.* **New CEOs.** *Indicator variable for CEO change in month.* **New Products.** *Indicator variable for new-product announcement in month.* **Earnings Announcement.** *Indicator variable for new-product announcement in month.* **New Patents.** *Indicator variable for patent grant in month.*

Trading-Strategy Returns by Holding Period

	$\hat{\alpha}$	$\hat{\beta}_{\text{Mkt}}$	$\hat{\beta}_{\text{SmB}}$	$\hat{\beta}_{\text{HmL}}$
1 Week	0.93 (0.13)	0.97 (0.01)	−0.13 (0.03)	0.14 (0.02)
2 Weeks	0.29 (0.08)	0.98 (0.01)	−0.10 (0.02)	0.16 (0.02)
3 Weeks	0.14 (0.07)	1.01 (0.01)	−0.13 (0.03)	0.19 (0.02)
4 Weeks	0.05 (0.06)	1.00 (0.01)	−0.06 (0.03)	0.15 (0.02)

Table 4: *Abnormal returns relative to the Fama and French (1993) factors of a zero-cost value-weighted portfolio that is long the quintile of stocks with largest short-run tilt in month* $\{m-1\}$ *and is short the quintile of stocks with the smallest short-run tilt in month* $\{m-1\}$ *for the first* $h \in \{1,2,3,4\}$ *weeks of month* m*. Regardless of the holding period, we report* $\hat{\alpha}$ *in units of % per month. Sample period: February 2002 to December 2010. Numbers in parentheses are standard errors. Reads: "Trading on the information in last month's short-run tilt generates positive abnormal returns during the first week, but these returns rapidly decay towards zero at longer holding periods."*

Abnormal Returns Controlling for Execution Risk

	$\hat{\alpha}$	$\hat{\beta}_{\text{Mkt}}$	$\hat{\beta}_{\text{SmB}}$	$\hat{\beta}_{\text{HmL}}$
Hi	0.11 (0.12)	0.98 (0.03)	−0.37 (0.05)	0.13 (0.05)
4	0.12 (0.13)	1.03 (0.03)	−0.21 (0.06)	0.08 (0.05)
3	0.13 (0.15)	1.08 (0.03)	−0.11 (0.07)	−0.03 (0.06)
2	−0.03 (0.17)	1.12 (0.04)	−0.03 (0.08)	−0.03 (0.07)
Lo	−0.36 (0.24)	1.08 (0.05)	0.14 (0.11)	0.05 (0.09)
Hi − Lo	0.47 (0.30)	−0.10 (0.07)	−0.51 (0.13)	0.08 (0.11)

Table 5: *Abnormal returns relative to the Fama and French (1993) factors of value-weighted portfolios sorted by residual short-run tilt controlling for execution risk as described in Equation 10. Execution risk is defined as the standard deviation of the implementation shortfall faced by institutional investors trading stock n in month m; see Equation 15. Data include period from January 2002 to December 2010. $\hat{\alpha}$ has units of % per month. Numbers in parentheses are standard errors. Reads: "Execution risk accounts for around half of the abnormal returns associated with cross-sectional variation in short-run tilt."*