The Path of Price Discovery: Real Time vs Trade Time*

Shrihari Santosh†

First version: April 2013
This version: May 2019

Abstract

I develop a dimensionless measure of the path of price discovery using public information shocks as an instrument for changes in firm value. I find that price discovery occurs largely through trading. Using high-frequency data, I find substantial cross-sectional variation in the speed of price of price discovery when measured in "real time." This heterogeneity, however, essentially disappears when measured in "trade time." Consistent with previous work work, I find that trade time is well approximated by the accumulation of transactions, not share volume, dollar volume, or turnover. These findings support the hypothesis of Kyle and Obizhaeva (2016b) that "information flows take place in the same business time as trading."

*I am grateful for comments and suggestions from Torben Andersen, Doug Diamond, John Cochrane, Julien Cujean, Valentin Haddad, Lars Hansen, Bryan Kelly, Serhiy Kozak, Damian Kozbur, Pete Kyle, Juhani Linnainmaa, Mark Loewenstein, Toby Moskowitz, Ľuboš Pástor, Amit Seru, Harald Uhlig, Wei Xiong, and seminar participants at the Arrowstreet, Chicago, Emory, Maryland, and Wharton. Any errors are my own. This paper was previously circulated as Public Information and Price Discovery.

†University of Colorado at Boulder shrihari.santosh@colorado.edu

Disclosure Statement

Shrihari Santosh declares that he has no conflicts of interest to disclose.

1 Introduction

Financial markets function as mechanisms for information aggregation as well as for capital allocation and risk sharing.[1] With some exceptions, the traditional view is that trading is not important for price discovery in response to public information.[2] This paper challenges that perspective. Using high-frequency data and a new approach for measuring price response to news, I find substantial cross-sectional variation in the speed of price of price discovery when measured in "real time." This heterogeneity, however, essentially disappears when measured in "trade time." Consistent with previous work work, I find that trade time is well approximated by the accumulation of transactions, not share volume, dollar volume, or turnover. Importantly, the *invariance* I document in trade time crucially depends on the natural rescaling inherent in my measure of price discovery.

I measure the path of price discovery by a cumulative impulse response function, θ_t, the change in expected stock price at various future "short horizons." Importantly, these are scaled by the shock in expected stock price at some "long horizon," essentially standardizing the shock to unity. If news is binary and completely observable by an econometrician, θ_t can be easily recovered as the ratio of regression coefficients when forecasting short and long horizon returns with news. Unfortunately, typical financial news is neither. I propose an instrumental variables methodology which resolves these two issues but requires a further assumption that conditional on some observable variables, the ratio is independent of the realized long run shock. I apply this procedure to quarterly earnings news, using analysts' earnings per share forecast error as an instrument.

I first measure θ_t in constant "real time" using intraday returns. Pooled estima-

[1]See Green (1973) for a model of information aggregation in asset markets. The basic idea that markets communicate information is developed by Hayek (1945), who notes that "we must look at the price system ... as a mechanism for communicating information if we want to understand its real function".

[2]Indeed, the microstructure literature typically assumes this process is instantaneous whereas the behavioral finance literature attributes documented predictability to the slow diffusion of information through non-trading mechanisms.

tion (ignoring heterogeneity) reveals a striking pattern. For earnings announcements which occur after the market closes, the next day opening price incorporates 74% of the news, indicating substantially price discovery takes place during extended hours trading. After the open, the dramatic increase in liquidity leads to rapid price discovery; θ_t increases to 81% within three minutes and is over 90% within one hour.[3] Pooled estimation, however, masks substantial heterogeneity. For the largest quintile of firms, θ_t is nearly 100% by the open whereas for the smallest quintile, it is only 68%. Then the small firms catch up; within thirty minutes the difference in θ_t shrinks to 11%. I split firms on various measures of liquidity and investor sophistication and find a similar pattern. There are large differences in θ_t at the open which largely disappear within the first 30-60 minutes of regular trading.

Motivated by literature linking the intensity of transactions with volume and volatility, I measure θ_t immediately after news arrival in constant "trade time." That is, I measure time as the number of transactions since the earnings announcement. Pooled estimation reveals that essentially all of price discovery is a "drift" with nearly zero jump in conditional mean prices as a result of news. This contrasts with a standard assumption in many microstructure models of instantaneous adjustment to public information.[4] However, price discovery proceeds rapidly; θ_t increases smoothly to 65% by the thirtieth transaction post-announcement. The median time elapsed since announcement to the thirtieth trade is only four minutes. Unlike in real time, θ_t measured in trade time exhibits little, if any, cross-sectional heterogeneity. This supports conjecture in Kyle and Obizhaeva (2016b) that "information flows take place in the same business time as trading." Further evidence that number of transactions is a better measure of business time than real time comes from comparing events with very different speed of trading. I sort announcements into two groups, *Fast* and *Slow,* based on the time elapsed from announcement to the thirtieth transaction. θ_t in trade time is nearly identical for the two groups, though they differ greatly in time elapsed; median time taken for *Fast* and *Slow* is one and fifteen minutes, respectively.

[3] Regular trading hours volume is an order of magnitude larger and bid/ask spreads are an order of magnitude smaller than during extended-hours trading (Barclay and Hendershott, 2004). Kelley and Tetlock (2013) exploit this variation in liquidity to estimate a structural model of overconfidence and trading.

[4] See for example, Copeland and Galai (1983); Roll (1984); Kelley and Tetlock (2013).

The trade time invariance I document provides a new approach understanding "business time," an idea that goes back at least to Mandelbrot and Taylor (1967). They introduce the idea that if transaction level returns are Gaussian, returns may be fat tailed (Pareto) when viewed over a fixed time period. Clark (1973) proposed measuring markets in cumulative volume time as a way of recovering normality of returns. Empirically, Jones et al. (1994) find that "average trade size has virtually no explanatory power when volatility is conditioned on the number of transactions." Motivated by previous work, Ané and Geman (2000) introduce the idea of stochastic time change and find that the distribution of returns conditional on the number of trades is well approximated by a normal distribution.

The findings in this paper also contribute to a large literature which measures the properties of price discovery. There are a number of works documenting the process in U.S. Treasury and foreign exchange markets. Fleming and Remolona (1999) find that excess volatility and volume in Treasuries persist for ninety minutes after monthly CPI, PPI and employment releases but claim that prices react "nearly instantaneous[ly]" to public information with no trading required. Andersen et al. (2007) find that stock, bond, and currency prices jump in response to news but they use five-minute returns and thus do not distinguish between true jumps and rapid "drift." In contrast, Green (2004) shows the asymmetric information component of bid/ask spreads on Treasury bonds increases after news releases, suggesting "that market participants watch trading to help determine the effect of new...information." Brandt and Kavajecz (2004) find that order flow imbalances result in permanent price changes even in the absence of macro news, evidence of learning through trading. Consistent with Kyle (1985), orders have larger price impact when liquidity is lower. Evans and Lyons (2002) find that daily movements in exchange rates and net order flow are strongly correlated. Evans and Lyons (2008) estimate that roughly two-thirds of the impact of news on exchange rates occurs through order flow. They propose that "dealers observe macro news but have little idea how to interpret it, or how the rest of the market will interpret it. Instead, they wait to observe the trades induced and set their prices and expectations based on the interpretations embedded therein". Their hypothesis is consistent with the evidence in this paper that price discovery, even for public information, occurs through trading.

The methodology in this paper is closest to Beechey and Wright (2009), who essentially run the reduced form version of the instrumental variables approach I employ. They find that ten-year Treasury and TIPS yields fully respond to most macroeconomic news releases within ten minutes. Other papers measure price discovery through excess volatility, volume, or bid-ask spreads. Any study of price discovery requires a model of the counterfactual: what would the measured quantity be in the absence of news? Though useful in some contexts, the previously employed methods are inappropriate for my setting. They all require that a single news shock occurs during the event measurement window. This may be a reasonable assumption for macroeconomic news but is likely violated for earnings announcements. Shortly after the earnings press release, firms typically hold a "conference call" in which the CEO/CFO gives more detail and soft information regarding the previous fiscal period. In addition, they often provide "guidance," or a forecast of future fundamental performance. A study of excess volatility might reasonably measure the time required for the entire "information episode" to be incorporated into prices, but cannot identify the price discovery of a single shock. Further, most stocks have very little extended hours trading on non-news days. This creates significant empirical difficulty in accurately measuring what volume, volatility, or spreads "would be" absent news. In contrast, the IV approach requires only an adjustment for expected returns, which are nearly zero over short horizons.[5] The cost of this robustness is that IV estimates may be less certain than these alternative measures.

2 Methodology

Let R_t be the cumulative return from just before an information shock (time 0) to event time t and R_{lr} be the "long-run" return some time after price discovery is complete.[6] Event studies in finance typically ask two questions: first, what is the average change in firm value given some news, $\Delta\text{Value} = \mathbb{E}\left[R_{lr} \mid \text{News}\right]$, and second,

[5] In other words, the "joint hypothesis problem" (Fama, 1970) is of little practical concern.

[6] I do not compute "abnormal" returns relative to an asset pricing model since over short horizons risk premia are essentially zero.

what is the path by which price reaches that value, $\mathbb{E}\left[R_t|\,\Delta\text{Value}\right]$, where R_t is the cumulative return from time 0 to t. When news is binary, the path is characterized by the "cumulative impulse response" ratio

$$\theta\left(t\right)=\frac{\mathbb{E}\left[R_t|\,\text{News}\right]-\mathbb{E}\left[R_t|\,\text{NoNews}\right]}{\mathbb{E}\left[R_{lr}|\,\text{News}\right]-\mathbb{E}\left[R_{lr}|\,\text{NoNews}\right]},\tag{1}$$

which may vary across events. Practically, most news is neither binary nor completely observable by an econometrician. Below I propose an event study methodology which addresses these issues, but restricts θ_t to depend only on observable variables. That is, conditional on observables, the path of price discovery is not allowed to depend on the realization of the news itself.

Let $\mathcal{I}_{i0}$ be an information set just prior to a news release for event i and assume price discovery is complete by some finite time T for all events. That is, the shock $\mathbb{E}\left[R_{it}|\,\mathcal{I}_{i0},\ \text{News}_i\right]-\mathbb{E}\left[R_{it}|\,\mathcal{I}_{i0}\right]$ is constant, ΔV_i, for all $t\geq T$. By iterated expectations, ΔV_i is conditionally mean zero. The key restriction which allows for continuous news is that:

$$\mathbb{E}\left[R_{it}|\,\mathcal{I}_{i0},\ \text{News}_i\right]-\mathbb{E}\left[R_{it}|\,\mathcal{I}_{i0}\right]=\theta\left(t,C_i\right)\cdot\Delta V_i,\tag{2}$$

where $C_i\in\{\mathcal{I}_{i0},\ \text{News}_i\}$ are K predetermined observable variables. The above implies that (conditional on C_i) a constant fraction, $\theta\left(t,C_i\right)$, of the "long-run" shock ΔV_i is expected to be incorporated into prices at horizon t. This linear form often obtains in models such as those Kyle (1985); Back and Pedersen (1998); Hong and Stein (1999) and in **??** I find little empirical evidence of non-linearity. Equation 2 does allow for correlation of ΔV and C, which implicitly allows for correlation of ΔV_i and θ_i, when *not* conditioning on C. For simplicity, I require a linear structure $\theta\left(t,C_i\right)=\delta_t'C_i$ where δ_t is an unknown parameter to be estimated.[7] For convenience, I use the notation $\theta_t\left(C\right)=\theta\left(t,C\right)$. Further, I require K instrumental variables, $Z\in\{\mathcal{I}_0,\ \text{News}\}$, which are sufficiently correlated with C_iR_{iT} and satisfy a weak orthogonality condition detailed in Appendix A. In practice I use as instruments s_iC_i, where s_i is a quarterly earnings surprise and C_i are variables chosen to generate heterogeneity in liquidity and investor sophistication. The instruments are necessary since ΔV is unobservable; we have only a noisy proxy R_T. If ΔV were observable, the

[7] It is possible to allow θ to be nonlinear in the parameters δ_t but requires estimation via nonlinear least squares.

linearity assumption, Equation 2, would not be required; a nonparametric regression of R_t on ΔV could recover the non-linear relationship.[8] Finally, one may also include controls $X \in \mathcal{I}_0$ (including a constant) which help to absorb $\mathbb{E}\left[R_{it}|\mathcal{I}_{i0}\right]$.

Given the above structure, $\theta_t(C)$ can be estimated by two stage least squares (2sls) as follows:

$$\text{1st stage}: E\left[C_i R_{iT}|Z, X\right] = \alpha' X + \beta' Z \tag{3}$$

$$\text{2nd stage}: \quad E\left[R_t|Z, X\right] = a_t' X + \delta_t' \widehat{E}\left[C_i R_{iT}|Z, X\right], \tag{4}$$

where E denotes linear projection (as opposed to $\mathbb{E}$ which denotes expectation). Note there there are as many 1st stage regressions as there are elements of C. When C includes a constant, it is convenient to partition δ_t' as $\left[\begin{array}{cc}\kappa_t & b_t'\end{array}\right]$ so that b_t measures the impact of the non-constant elements of C on θ_t.

2.1 Invariance

The idea of time change and invariance goes back at least to Mandelbrot and Taylor (1967) and Clark (1973) who study time change and normality of returns. Jones et al. (1994) find that cross-sectionally "average trade size has virtually no explanatory power when volatility is conditioned on the number of transactions." Motivated by previous work, Ane and Geman find that, in fact, the distribution of returns conditional on the number of trades is very well approximated by a normal distribution. Kyle and Obizhaeva (2016b) develop market microstructure invariance hypotheses regarding the relationship between business time, risk transfer, and transaction costs, but "conjecture that data on news articles can help to show that information flows take place in the same business time as trading." Using dimensional analysis to relate volume, volatility, trade frequency, and trade size, Kyle and Obizhaeva (2016a) find empirically that number of trades per unit time (across stocks) proxies well for the number of unobserved "bets" (unique trading decisions). Taken together, there is strong support for the idea that equilibrium "business time" is well measured by the cumulative number of transactions rather than other measures such as turnover,

[8] Recent methods for non-linear instrumental variables with nonclassical measurement error such as that proposed by Hu and Schennach (2008) are not applicable here since the error in R_T is likely correlated with the error in R_t.

or share or dollar volume. In the context of price discovery after public news, the "right" measure of business time, or "clock," is one which produces invariance, or *lack* of cross-sectional heterogeneity in θ. That is, $\theta(t, C)$ collapses to $\theta(t)$ when t is appropriately measured.

It is, however, not obvious that such a clock even exists. To maintain irreversability of time, we must require that even if the passage of business time is stochastic (as in Ané and Geman (2000)), ex post measured business time elapsed must be a monotone transformation of real time elapsed. Now suppose that in real time, $\theta(t, C_i)$ is monotone in t for some C_i but nonmonotone for some C_j. Then there is *no* clock, τ, which generates invariance, $\theta_i(\tau) = \theta_j(\tau)$ for all τ since $\theta_i(\tau)$ is just a nonlinear (but monotone) dialation of $\theta_j(t)$. The empirical estimates in **??** and Section 4 suggest that θ exhibits high-frequency invariance when measured "trade time" around the announcement, but displays very clear calendar dependence due to real time effects such as the opening and closing of markets.

2.2 Event Timing

Unlike previous studies, I rely on neither IBES nor Compustat for earnings announcement dates since these are known to contain date and time errors. From Wall Street Horizon (WSH), I obtain earnings press release dates and times, recorded to the nearest second.[9] WSH is a for-profit company exclusively devoted to accurate recording and prediction of earnings announcement, conference calls, and ex-dividend dates and times as well as actual and estimated EPS values. The sample includes announcements for all US domiciled public firms from 2006 to 2011. Figure 1 below shows the distribution of earnings announcement times (within day). The bimodal distribution has little mass during trading hours. In the main analysis I restrict to announcements after 16:00 (after market close). In Section 5.3 I show that price discovery of announcements before 9:30 (before market open) is qualitatively similar.

I measure cumulative event returns starting at the last closing price *before* the announcement. I refer to this event time as C0. After-hours trading continues (potentially) until 20:00. Pre-market trading begin (potentially) at 4:00 the next day.

[9] I manually verify 200 announcement times using Factiva and find 100% agreement between WSH data and the press release timestamp.

Figure 1: Histogram of Earnings Announcement Times

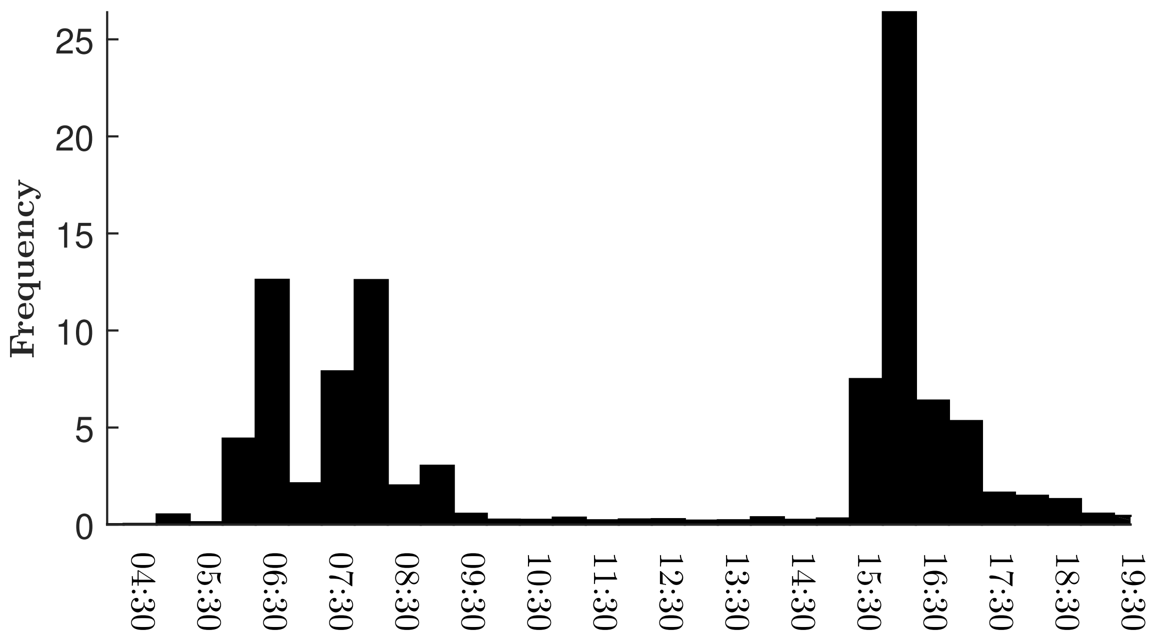

30 minute intervals. Time shown is (inclusive) left endpoint.

Extended hours trading takes place on NASDAQ, NYSE ARCA, and various other ECNs (electronic communication networks). Mechanically it is very similar to regular trading on the NASDAQ, where traders may both supply liquidity by posting limit orders or consume liquidity by executing immediately against an open order. A major regulatory difference is that the SEC's "Order Protection Rule" (Reg. NMS, Rule 611) does not apply during extended hours. Rule 611 integrates all exchanges during regular trading, essentially requiring that orders execute at the national best bid or best ask (best quote across exchanges). Practically, extended hours trading is very different from the regular trading session. Extended hours bid/ask spreads are at least ten times larger and trading volume twenty times smaller relative to regular trading.[10] As shown in Figure 4a, not all firms trade after hours, even after an earnings announcement. The first time all events trade is 9:30 (open) the next day. I refer to the first opening price post announcement as O1. Similarly, C1 represents the

[10]See Barclay and Hendershott (2003, 2004, 2008) for in depth description and analysis of extended hours trading.

closing price *after* one full regular trading session. Subsequent opening and closing prices are numbered sequentially.[11]

2.3 Trade Data

I obtain opening and closing prices for individual stocks from CRSP.[12] NYSE and NASDAQ have similar auction mechanisms to set the official opening and closing prices.[13] θ estimated separately on NYSE and NASDAQ are quantitatively similar, indicating that any differences in trading mechanism do not substantially affect price discovery. Intraday trade data is from NYSE TAQ, which includes all trades and quotes reported on the consolidated tape.[14] For pre- and after-hours trades, TAQ includes all transactions of at least 100 shares, but does not include trades on exchanges which do not participate in the National Market System. Barclay and Hendershott (2008) show that TAQ "is a reliable source for after-hours trades" though it technically does not contain the complete trading record.

In Section 3 I measure θ_t in constant "real time" during regular trading. Whereas nearly all equities (at least those which are covered by an institutional equity analyst) trade each day, the intraday trading record is more sparse. Therefore higher frequency of observation allows for finer analysis but results in more missing data, especially for smaller firms. Balancing these two issues, I construct three minute volume-weighted average prices (VWAP) from 9:30 to 16:00. This still results in many missing prices, especially for smaller firms. To address this issue, I use a Brownian bridge for random imputation.[15]

[11] I use this convention for all announcements, whether they occur after market close or before market open.

[12] I include only firms listed on NYSE, Amex, NASDAQ, and NYSE Arca exchanges and securities with CRSP share codes 10 and 11 (Ordinary Common Shares with no further designation). This excludes ADRs, ETFs, REITs, and closed-end mutual funds.

[13] See https://www.nyse.com/publicdocs/nyse/markets/nyse/NYSE_Opening_and_Closing_-Auctions_Fact_Sheet.pdf and https://www.nasdaqtrader.com/trader.aspx?id=openclose

[14] Following Huang and Stoll (1996); Bessembinder (1999), during regular market hours I include trades with *Sale Condition* of Blank, @, F, E, @F, and @E. This process excludes trades that are out of sequence, involve error corrections or non-standard settlement. For extended hours, I also include code T which indicates NASDAQ extended session trades. I further exclude transactions with zero price or volume. I exclude block transaction of >10,000 shares since they are likely to be used by institutions for liquidity trading.

[15] To interpolate missing prices, I assume log prices follow a (potentially heteroskedastic) diffusion

In Section 4 I measure θ_t in "trade time" (trade-by-trade) in close proximity to the announcement. This leads to a different type of missing data problem since trading for many firms ceases within a few transactions post-announcement. Since trading (during the relevant session) ceases completely, it is not possible to interpolate using "nearby" prices. In Section 4 I show that missing trades do not result in any discernible bias.

2.4 Earnings Surprises

I construct realized earnings surprises using the Thomson-Reuters IBES database. For each quarterly earnings announcement by a U.S. firm, I construct $\widehat{eps}$ as the mean of analysts' earning per share forecasts in the inclusion window. I follow Thomson-Reuters' convention and define the window as follows: for fiscal quarters 1, 2 and 3, the inclusion window starts 105 calendar days prior to the earnings announcement. For fiscal quarter 4, the inclusion window starts 120 days prior, since Q4 announcements typically occur more days after the end of the fiscal period. Since most estimates are revised shortly after an announcement, this extended window allows extra time for companies to report year-end results.

If an analyst issued multiple forecasts in the inclusion window, I use only the most recent one. I exclude announcements which occur more than 80 days after the end of the fiscal quarter (as recorded in IBES) and limit the sample to announcements that have valid stock return data in CRSP continuously for at least 20 trading days prior to the press release date. Furthermore, I exclude earnings announcements that are either less then 45 days or more than 150 days since the previous earnings announcement from the same company. This procedure excludes earnings events where the company did not report earnings the previous quarter or where the fiscal calendar changed, resulting in a significantly shortened quarter.

with zero drift. Given two observations from a diffusion with known variance, the conditional normal distribution of intermediate values is given analytically (this is known as a Brownian bridge when the process is homoskedastic). Conditional means are derived from linear interpolation on the known values and conditional variances follow a quadratic, with zeros at the two known points. Since I do not know the conditional variance of the process, I use estimates obtained from the cross-section of non-missing returns. The procedure draws random values from the distributions given by the analytic formulas. To avoid extrapolation based on the imputation assumptions, I require non-missing opening and closing prices.

Table 1: Horizon Choice

$$\text{1st stage}: E\left[R_T|\, s\right] = \alpha + \beta s$$
$$\text{2nd stage}: \widehat{E}\left[R_t|\, s\right] = c_t + \theta_t \cdot \widehat{E}\left[R_T|\, s\right]$$

Each column gives the estimate of θ_t at the indicated horizon. The "long-run" horizon, T, is ten trading days post-announcement. Standard errors (in parentheses) are two-way clustered by firm and quarter of announcement.

	O1	C1	C2	C3	C5	C7	C9
θ_t	73.6	95.7	100.4	100.4	101.6	100.5	100.5
	(3.2)	(3.6)	(3.6)	(3.5)	(3.2)	(2.6)	(2.3)

Following Kothari (2001); DellaVigna and Pollet (2009), I define the normalized earnings surprise as the difference between actual and forecast earnings per share (EPS), scaled by lagged stock price. Letting $eps_{i,q}$ and $\widehat{eps}_{i,q}$ represent the actual and forecast earnings per share for company i in quarter q and letting $\bar{P}_{i,q}$ be the split-adjusted closing price per share of company $i's$ stock five trading days prior to the earnings announcement for quarter q, I define the normalized earnings surprise $s_{i,q}$ as:

$$s_{i,q} = \frac{eps_{i,q} - \widehat{eps}_{i,q}}{\bar{P}_{i,q}}. \tag{5}$$

I exclude observations where $\|s_{i,q}\| > 0.01$ (which trims the data at the 5th and 95th percentiles). Finally, I exclude observations for which lagged price is below \$5/share. Each quarter, the included firms represent approximately 70% of the total market value of the CRSP universe. For notational convenience I drop q subscripts so s_i refers to the earnings surprise for event i (which implicitly refers to a firm,quarter).

3 Real Time

For "pooled" estimation ($C_i = 1$) I use the earnings surprise s_i as the single instrument, but results are robust to inclusion of higher order terms (i.e. s_i^2 and s_i^3), or sorting s_i into bins and using a vector of indicator variables (see **??**). This still leaves

Table 2: Event Day Impulse Response

Each column displays the change in θ_t (in percent) measured over the given horizon. For example, the third column reports 7.87 for the period 10:00-3:45, meaning the cumulative impulse response increases by 7.87% during this time frame. Standard errors are adjusted for heteroskedasticity and are clustered by firm and quarter of announcement. Absolute t-statistics in parentheses.

	O1→9:33	9:33→10:00	10:00→3:45	3:45→C1	C1→O2	O2→C2
$\theta_{\omega:t}$	6.70**	8.51**	7.87**	-1.22**	0.93	3.85**
	(8.0)	(7.5)	(6.8)	(4.3)	(1.9)	(4.3)
N	13535	13535	13535	13535	13535	13535
# Firms	1627	1627	1627	1627	1627	1627
# Year× Quarter	24	24	24	24	24	24

* $p < 0.05$, ** $p < 0.01$

the "long-run" horizon event date unspecified. The choice is not innocuous since by assumption all relevant information from the event is impounded into prices by this time. The framework proposed in Section 2 implies choosing any $\tilde{T}$ which is "long enough" produces consistent estimates of θ_t, but the uncertainty of those estimates increases with $\tilde{T}$.[16] Supposing we know that price discovery is complete by some time $\tilde{T}$, but may be complete sooner. Then we can use $\tilde{T}$ as a provisional long-run, then find the shortest $t^\star$ such that $\theta_t = 100\%$ for all $t \geq t^\star$. We then choose $T = t^\star$ for further analysis. Table 1 presents estimates of θ_t assuming twenty trading days is sufficient for price discovery to be complete.[17] The estimtates show that $\theta_t \approx 100\%$ at two trading days post announcement (C2), and remains so at all longer horizons. Given this finding, I use $T = C2$ in the analysis below. Using any shorter horizon produces estimates which are biased upward. Using any longer horizon results in less precise (though consistent) estimates of θ_t.[18]

The estimates in Table 1 imply there is no measurable "drift" after two trading days post-event. Lack of drift after two days contrasts with prior studies which find significant predictability for months after earnings announcements. The difference could come from sample selection since I exclude extreme surprises. It may also

[16] This is due additional "noise" which cumulates with time.

[17] Standard errors are two-way clustered by firm and quarter of announcement (see Thompson, 2011; Cameron et al., 2012).

[18] For before market open announcements presented in **??**, I use $T = C5$ based on analysis similar to that in Table 1.

Figure 2: real Time Cumulative IRF

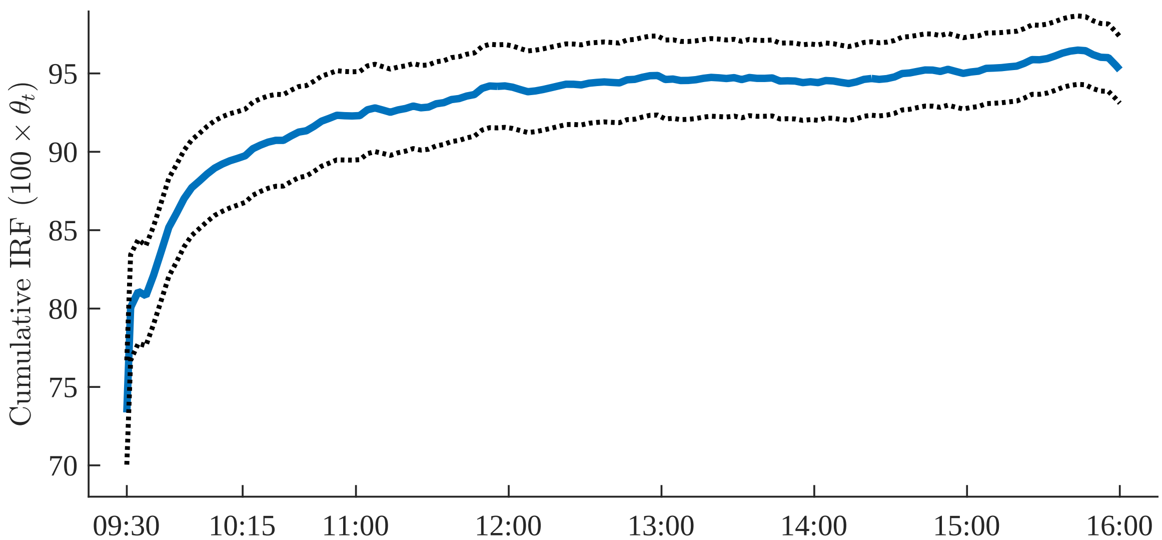

Trade 0 is the last trade prior to the earnings press release. Trade 1 is the first trade post announcement. Dotted black lines show two standard error bounds.

arise from data timing; I estimate using earnings from 2006-2011 whereas earlier studies typically use data from 1990-2005. I apply the methodology of DellaVigna and Pollet (2009) using data from 2006-2011 and find drift for at least one month. After trimming the data at the 5th and 9th percentiles of the surprise distribution I find predictability for only two days, consistent with Table 1. It seems that in a more modern sample longer horizon drift is a phenomenon isolated to extreme surprises, which tend to happen only for very small firms.

3.1 Pooled Estimation

Since many stocks have little to no extended hours trading even after an earnings announcement (see Figure 4a), I measure price discovery in "real time" starting with the next day regular market open, O1. Figure 2 plots pooled estimates of the cumulative impulse response, θ_t, starting at O1 continuing to the following close, C1. The leftmost point, θ_{O1}, shows that nearly 75% of price discovery is complete by the opening auction, suggesting that (on average) extended hours trading may be important. The figure shows rapid adjustment during the first hour of regular trading, especially in the first few minutes. Importantly, this pattern at the open of regular trading only occurs for announcements with limited after-hours trading (see Figure 5b below). The rightmost point, θ_{C1}, indicates that 95% of price discovery is complete after one full trading day.

Table 2 presents estimates of $\theta_{\omega:t} \equiv \theta_t - \theta_\omega$ over time intervals. For example, the first column reports 6.7 which equals $\theta_{9:33} - \theta_{\mathrm{O1}}$. It is the change in cumulative price reaction from the opening auction price to the 9:30-9:33 VWAP. The next two columns shows that price adjustment to news continues throughout the entire trading day in an economically and statistically significant way. The anomaly is the fourth column, which measures price reaction during the closing fifteen minutes of trading. The non-monotonicity observed in Figure 2 is statistically significant, though small. This retracing of prices is possibly due to a discontinuity in trading costs at 16:00. NASD Rule 2520 allows day traders to employ up to 400% leverage for intraday equity positions whereas positions held overnight are subject to the usual 200% leverage restriction. This rule greatly increases the cost of holding a position after 16:00. This reversal is reminiscent to the finding in Cushing and Madhavan (2000) that increased demand for immediacy near the close generates return reversals. The last two columns measure price discovery over the next day. They show that of the 5% remaining after the first day, nearly all of price discovery occurs during regular trading hours.

3.2 Heterogeneity

Table 3: θ_t by Liquidity/Competition Measure

Each column measures the cumulative impulse response (in %) by proxy at the indicated horizon. $\theta_\text{t} = \kappa_t + b_t' \mathbf{1}_{\text{Liq/Comp}}$. Standard errors are adjusted for heteroskedasticity. Absolute t statistics in parentheses.

	O1	9:30	10:00	10:30	12:30	C1
κ	97.4	97.5	96.4	97.8	102.7	97.0
	(15.6)	(20.7)	(23.1)	(22.6)	(27.2)	(23.6)
b_1	−29.3**	−21.7**	−11.0*	−11.2*	−11.1*	−5.7
	(4.3)	(4.2)	(2.2)	(2.0)	(2.2)	(1.1)
b_2	−23.2**	−15.9**	−6.9	−4.8	−5.3	2.6
	(3.6)	(2.8)	(1.6)	(1.2)	(1.2)	(0.5)
b_3	−18.7**	−15.2**	−5.6	−3.0	−8.9	−0.6
	(2.8)	(2.8)	(1.0)	(0.5)	(1.6)	(0.1)
b_4	−9.4	−6.0	1.9	1.0	−3.8	3.7
	(1.2)	(1.0)	(0.4)	(0.2)	(0.7)	(0.7)
κ	81.6	87.7	94.7	97.6	99.2	98.5
	(27.9)	(30.3)	(33.4)	(34.4)	(47.3)	(59.5)
$b_{\text{HighRetail}}$	−15.0**	−14.0**	−11.2**	−12.5**	−8.8**	−6.0**
	(5.0)	(4.5)	(3.3)	(3.5)	(3.0)	(2.9)
κ	78.8	84.1	91.0	93.0	96.1	97.6
	(34.0)	(35.7)	(45.0)	(45.2)	(65.6)	(77.1)
$b_{\text{LowCoverage}}$	−11.4**	−8.4*	−5.0	−4.7	−3.6	−5.0**
	(3.2)	(2.2)	(1.5)	(1.4)	(1.6)	(2.8)
κ	81.1	85.0	90.3	93.5	96.3	97.9
	(33.4)	(33.4)	(38.7)	(38.5)	(57.8)	(66.4)
$b_{\text{HighAmihud}}$	−12.4**	−7.9*	−2.8	−4.5	−3.1	−4.2*
	(3.0)	(2.0)	(0.8)	(1.4)	(1.2)	(2.3)
κ	82.7	86.3	92.4	95.4	97.1	98.3
	(76.3)	(53.5)	(93.6)	(57.7)	(46.0)	(62.5)
$b_{\text{HighSpread}}$	−12.8**	−8.5**	−5.2*	−6.4*	−3.7	−4.2*
	(5.0)	(3.0)	(2.1)	(2.5)	(1.3)	(2.0)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

The pooled estimates above recover the (cross-sectional) average path of price discovery, potentially ignoring substantial heterogeneity. It is plausible that price discovery may be slower for firms with illiquid stocks or fewer sophisticated investors. I explore such cross-sectional variation using various measures of liquidity and informed competition. First, I use holdings data from quarterly 13F filings to sort firms into two groups each calendar quarter based on (lagged) institutional ownership and set HighRetail = 1 if institutional ownership (as a percentage of shares outstanding) is below median. Since active traders are the clients of equity analysts, the number of analysts may be a good proxy for the level of competition among active investors. Further, more analyst coverage is thought to reduce information asymmetry through broader distribution of information. I define an event as having LowCoverage = 1 if only one analyst provided an earnings estimate for that quarter.[19] I also include three measures of liquidity. I calculate the price impact measure from Amihud (2002) using daily data in the prior calendar year and set HighAmihud = 1 if it is above median. Bid/ask spread is another commonly used measure of illiquidity based on the theory in Glosten and Milgrom (1985). I calculate average daily closing percentage spread, $\mathbb{E}\left[\frac{1}{2}\frac{ask-bid}{ask+bid}\right]$, for each firm in the prior calendar year and set HighSpread = 1 if it is above median. Previous studies have concluded that post-earnings announcement drift is a phenomenon concentrated in small firms (Bernard and Thomas, 1990; Chan, 2003; Hou, 2007; Vega, 2006; Chordia et al., 2009). This is rationalized by the observations that small firms have a higher percentage of (potentially biased) retail investors and higher trading costs which may deter informed arbitrageurs from trading aggressively. In addition to the aforementioned liquidity and competition measures, I use market equity quintiles since firm size is often used as a "catch all" for unobserved but important frictions.

Using the parameterization $\theta_t = \kappa_t + b_t\,[\text{Measure} = 1]$, where Measure is one of the above variables, the effect of these variables on the speed of price discovery can be estimated using the interacted two stage least squares setup given in Equation 4. As instruments I use s_i and $(\text{Measure} = 1)\,s_i$. The parameter κ_t is the cumulative impulse response for the baseline group (Measure = 0) and b_t captures the difference between the Measure = 1 and Measure = 0 groups. For size quintiles, κ is the impulse response of the largest quintile of firms and each $b_i,\ i = 1 \ldots 4$ is the difference in θ between firms in the indicated size quintile and the largest quintile.

Table 3 shows the estimates from O1 to C1. The estimated b for each measure

[19] LowCoverage = 1 for approximately 37% of events.

Figure 3: Trade Time Cumulative IRF

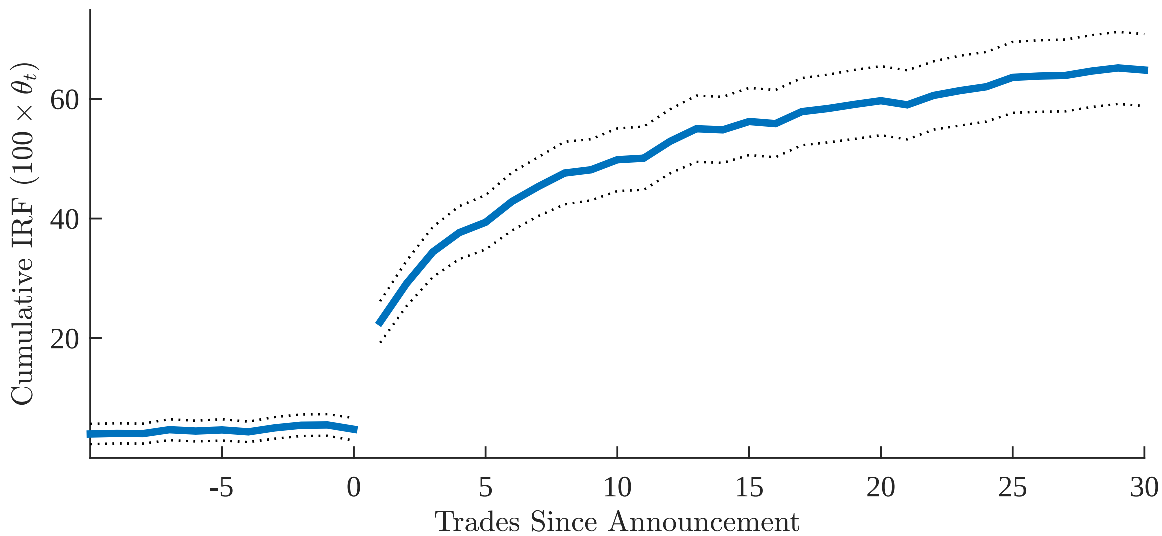

Trade -1 is the last trade prior to the earnings press release. Trade 0 is the first trade post announcement. Dotted black lines show two standard error bounds.

shows a large spread at the open which substantially decays over the trading day, particularly during the first hour. For the largest firms the opening price reflects essentially 100% of the earnings news ($\kappa = 97.8\%$). This is consistent with the findings of Gerlach and Lee (2011), who study stocks in the DJIA and measure speed of adjustment using excess volatility. For smaller firms there is rapid, but incomplete convergence within the first thirty minutes of trading. The other cross-sectional measures generate similar results. Cumulative price discover at the open is substantially lower for firms with low liquidity, analyst coverage, or high retail ownership. Measured in calendar time, there is large cross-sectional variation in the speed of price discovery.

4 Trade Time

The previous section shows that measured in real time, there is large cross-sectional variation in the speed of price discovery. Motivated by the previous research linking trades with volatility, I consider an alternative "clock." In this section I measure time

Figure 4: Trade Characteristics

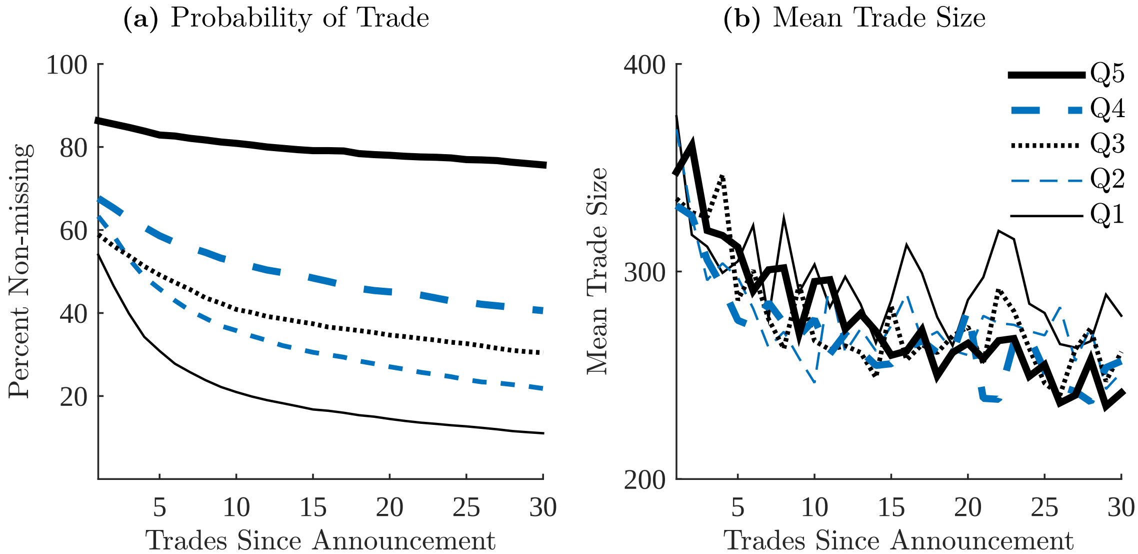

The left panel plots the the fraction of announcements which have at least the indicated number of trades post-announcement (by market equity quintile). The right panel plots 5% trimmed mean trade size. Trade 0 is the last trade prior to the earnings press release. Trade 1 is the first trade post-announcement.

as number of trades since the announcement. Because the data from WSH is time-stamped only to the nearest second, I exclude trades that are within ±1 second of the earnings announcement. Results are essentially unchanged if I include all trades. I require at least thirty transactions between the announcement and 20:00 (when the consolidated tape closes).

Figure 3 shows the estimated impulse response function around the announcement. Impulse response greater than zero at trade 0 is either due to information leaking after the close of regular trading but prior to the announcement or due to a potential bias in my timing data.[20] The most striking feature of Figure 3 is the small jump from trade 0 to trade 1. θ jumps to only 22% but increases rapidly to 65% by the thirtieth trade (median time of 4.5 minutes post-announcement). This pattern contrasts strongly with assumptions in many microstructure models of instantaneous adjustment to public information (Copeland and Galai, 1983; Roll, 1984, among others). Essentially

[20]WSH records when the earnings press release comes over the newswires. Occasionally, a company posts the information on its website just prior to the newswire release.

Figure 5: Selection Bias

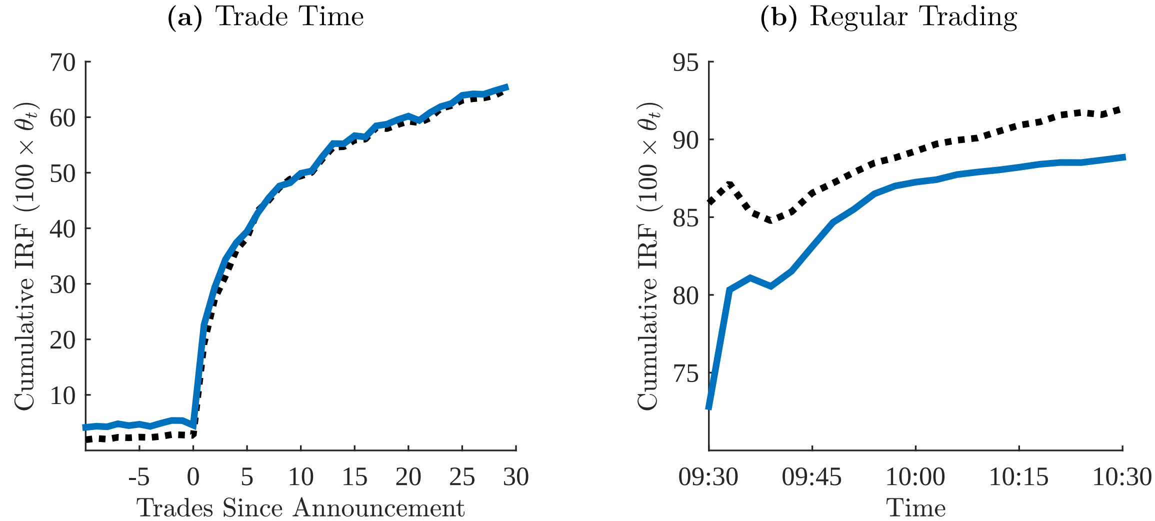

Trade 0 is the last trade prior to the earnings release. Trade 1 is the first trade post announcement. The solid blue function is estimated from all events with a valid trade at the indicated horizon. The dotted black function is estimated using only events with at least thirty post-announcement transactions, as in Figure 3.

all of price discovery is a "drift" with nearly zero jump in conditional mean prices as a result of news. This finding is consistent with a model featuring strategic traders and "dealers [who] observe ... news but have little idea how to interpret it, or how the rest of the market will interpret it. Instead, they wait to observe the trades induced and set their prices and expectations based on the interpretations embedded therein" (Evans and Lyons, 2008). The reader may be concerned that slow adjustment is due to "stale" limit orders from inattentive retail traders. Most retail brokerages, however, either do not allow after hours trading or default to regular trading hours execution for good-till-canceled orders.

The above analysis may be subject to ex-post selection bias since I only include announcements with at least thirty after hours trades. Indeed, Figure 4a shows that such the requirement biases my sample towards larger firms. The plot shows the fraction of announcements which have at least the indicated number of trades post-announcement. There is a extreme cross-sectional heterogeneity: nearly 80% of large firm, but fewer than 20% of small firm announcements are accompanied by at least

Table 4: Summary Statistics by Trading Speed

Mean, median, and standard deviation for indicated measure by trading speed.

	Mean		Median		σ	
	Fast	Slow	Fast	Slow	Fast	Slow
Time to 30th Trade (Minutes)	2.1	18.7	1.0	15.0	2.7	15.2
Surprise (%)	0.1	0.1	0.1	0.1	0.2	0.2
R_{C2} (%)	−0.1	−0.1	0.1	−0.2	10.7	9.9
R_{T30} (%)	−0.1	−0.5	0.1	−0.0	4.8	5.6
Inst. Ownership	0.8	0.8	0.8	0.8	0.2	0.2
Analyst Rec.	2.4	2.4	2.4	2.4	0.6	0.6
ME Quintile	3.3	3.3	3.0	3.0	1.4	1.4
BM Quintile	2.2	2.4	2.0	2.0	1.2	1.3

thirty after-hours trades after the earnings press release. This heterogeneous inclusion and the large heterogeneity in speed documented in Table 3 suggests that selection bias may be large.

I re-estimate the cumulative impulse response at each horizon using all events with a non-missing price to assess the extent of potential bias. Figure 5a shows this function (dotted black) along with the estimates from Figure 3 (solid blue). The two functions essentially overlap, indicating that estimates using only observations with non-missing data capture the essential features of the impulse response. Figure 5b shows the impulse response in the first hour of regular session trading separately for two groups: those observations with a valid thirtieth trade (dotted black) and those without (solid blue). Firms with non-missing after-hours prices have significantly higher price discovery at the open but within one hour the difference is largely eliminated as the remaining firms catch up. This finding mimicks the result in Table 3 and futher highlights the importance of extended hours trading for price discovery

4.1 Real Time vs Trade Time

If price discovery is invariant in trade time, the impulse response function should only depend on the cumulative number of transactions and *not* on the real time elapsed. To test this hypothesis, I sort firms based on the real time elapsed from the earnings release to the thirtieth trade into two equal size groups (Fast and Slow), within each size quintile, within each calendar quarter (Year × Quarter). This method orthogonalizes speed relative to firm size and calendar time which is important since

Figure 6: Cumulative IRF by Trading Speed

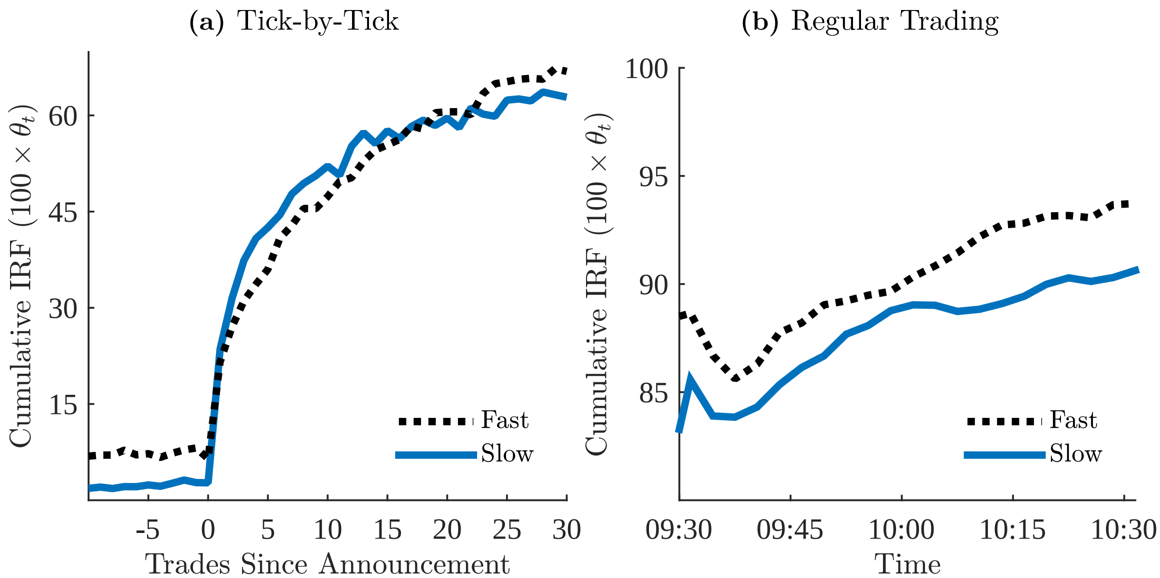

Trade - is the last trade prior to the earnings press release. Trade 1 is the first trade post announcement. Firms are sorted by duration from announcement to thirtieth trade, within year, quarter, and size quintile.

speed is strongly negatively correlated with firm size. Median elapsed time for Fast and Slow firms are one and fifteen minutes, respectively. Table 4 shows that Fast and Slow firms and indistinguishable on a variety of other characteristics. This mitigates concerns of bias since speed is defined based on ex-post elapsed time. Figure 6a shows θ_t (Fast) and θ_t (Slow) around the earnings announcement; the functions are nearly identical. A formal test (not shown) does not reject equality. This shows that, at least when measured around the announcement, the cumulative number of transactions is better than elapsed real time as a measure of "business time".

4.2 Heterogeneity

Table 5: CIRF_t by Liquidity/Competition Measure

Each column measures the cumulative impulse response (in %) by proxy at the indicated horizon. $\theta_t = \kappa_t + b_t\,(\text{Measure} = 1)$. The last column gives estimates (in %) from $\text{ElapsedTime} = \kappa_t + b_t\,(\text{Measure} = 1)$, where ElapsedTime is the time from announcement to the thirtieth trade. Standard errors are clustered by firm and quarter of announcement. Absolute t statistics in parentheses.

	Trd_0	Trd_5	Trd_{10}	Trd_{20}	Trd_{30}	Time
κ	10.4	42.6	49.3	56.2	66.2	8.1
	(2.7)	(5.7)	(7.4)	(9.1)	(10.5)	(33.4)
b_1	−6.1	−0.2	1.2	7.0	−0.4	11.6**
	(1.4)	(0.0)	(0.1)	(0.9)	(0.0)	(15.5)
b_2	−7.9	−2.1	3.8	9.7	3.2	9.8**
	(1.9)	(0.2)	(0.4)	(1.0)	(0.3)	(16.0)
b_3	−5.1	−9.1	−1.9	−2.9	−4.4	8.4**
	(1.2)	(1.1)	(0.3)	(0.4)	(0.7)	(14.1)
b_4	−6.2	−4.1	−0.4	2.1	−4.1	5.4**
	(1.4)	(0.5)	(0.0)	(0.3)	(0.5)	(9.7)
κ	5.0	37.7	48.9	60.5	66.3	
	(3.1)	(12.3)	(14.4)	(14.9)	(16.5)	
$b_{\text{HighRetail}}$	−0.8	3.4	2.1	−0.7	−3.0	−2.7**
	(0.4)	(0.7)	(0.4)	(0.1)	(0.4)	(6.0)
κ	5.6	37.7	48.3	58.2	63.7	
	(3.2)	(10.3)	(12.2)	(13.5)	(15.1)	
$b_{\text{LowCoverage}}$	−3.8*	5.8	5.4	6.7	4.0	6.9**
	(2.1)	(1.1)	(1.0)	(1.0)	(0.6)	(10.4)
κ	5.0	35.5	46.4	56.7	63.3	
	(2.8)	(9.6)	(11.9)	(12.7)	(16.0)	
$b_{\text{HighAmihud}}$	−1.4	11.3*	10.3	10.5	5.2	7.4**
	(0.6)	(2.2)	(1.5)	(1.6)	(0.9)	(12.3)
κ	4.7	34.3	44.5	52.7	59.4	
	(3.0)	(9.4)	(12.1)	(11.9)	(15.3)	
$b_{\text{HighSpread}}$	−0.2	8.9	9.5	13.6*	10.2	8.0**
	(0.1)	(1.8)	(1.6)	(2.3)	(1.7)	(15.6)

* $p < 0.05$, ** $p < 0.01$

In Table 3 I found that in calendar time, there is large variation in the speed of price discovery which is correlated with common proxies of liquidity and investor sophistication. If price discovery is invariant in trade time, these differences should disappear when measured using that clock. To test this, I repeat the estimation of the cross-sectional specification $\theta_t = \kappa_t + b_t\,[\text{Measure} = 1]$, except that t is trade time. Table 5 shows that indeed, the differences largely disappear. The last column gives estimates from the regression $\text{ElapsedTime} = \kappa_t + b_t\,[\text{Measure} = 1]$, where ElapsedTime is the time from announcement to the thirtieth trade. The estimates show that firms with illiquid stocks or fewer sophisticated investors tend to trade much slower. These findings confirm that the invariance observed in Figure 6a is not an artifact of sorting on ex-post elapsed time.

Table 6 shows various summary statistics by market equity quintile. The last row gives the standard deviation of the first stage fitted value when R_{C2} is regressed on the earnings surprise (separately by quintile). This provides a lower bound on the standard deviation of the fundamental innovation, ΔV, and measures the "typical" size of a shock. There is a clear monotonic pattern; the smallest quintile of firms experience shocks which are three time as large as those experienced by the largest quintile. This difference highlights the importance in measuring price discovery with the dimensionless quantity θ. If one measures price discovery at time t as $std\left[\hat{\mathbb{E}}\left(R_t|\, S\right)\right]$ (the standard deviation of predicted returns), it will appear that small firms have much faster price discovery. If instead one measures the amount of "drift" or delayed reaction, $std\left[\hat{\mathbb{E}}\left(R_T - R_t|\, S\right)\right]$, it will appear that small firms have much more slower price discovery since their longer horizon returns display greater predictability. Invariance in the speed of price discovery only obtains when appropriately scaling according to Equation 2. Table 6 also shows that though price discovery may occur through trading, the process does not require substantial trading volume when markets are illiquid. Barclay and Hendershott (2003) come to a similar conclusion by estimating the structural model of Easley et al. (1996, 1997b,a).

4.3 Alternative Clocks

The above findings show that cumulative number of transactions is better than elapsed real time as a measure of "business time," but it is possible that some alternative may be better still. I consider three plausible alternatives: cumulative turnover,

Table 6: Trading Statistics by Firm Size

The first three rows are median values by market equity quintile for various measures over the first thirty trades post-announcement. Rows four is the standard deviation of predicted returns at indicated horizons.

	Small	2	3	4	Big
MarketCap. ($ B)	0.40	1.01	1.99	5.10	20.09
Volume ($ 100K)	1.72	2.41	2.86	3.72	3.79
Turnover (b.p.)	4.32	2.47	1.46	0.71	0.17
$std\left[\widehat{\mathbb{E}}\left(R_{C2}\middle\| S\right)\right]$ (%)	5.27	3.88	3.47	2.42	1.67

dollar volume, and share volume. Turnover and dollar volume can be ruled out by examination of Table 6. The third row shows that turnover of the smallest firms is more than twenty times the turnover of the largest firms, but Table 5 shows that speed measured in trade time is indistinguishable between these groups. Similarly, dollar volume of the largest firms is more than twice that of the smallest firms, with no difference in speed.

However, no such comparison can rule out share volume. Figure 4b shows (5% trimmed) mean trade size by market equity quintile. Remarkably, there is essentially no variation in average trade size. Therefore, the trade time invariance shown Table 5 is consistent with cumulative share volume being the right measure of "business time." To test whether price discovery is invariant when time is measured as cumulative share volume, I first sort events into two groups, Fast and Slow, based on the number of trades required for cumulative volume to reach 3000 shares. If the impulse response function for these two groups is significantly different, we can reject volume time invariance.

A practical difficulty with volume time is that it generates an unbalanced panel since trades can be of arbitrary size. As a solution, I impose that $\theta(t)$ is a continuous function of volume time, with a possible discontinuity at $t = 0$. I further restrict the

Figure 7: Cumulative IRF by Trading Speed

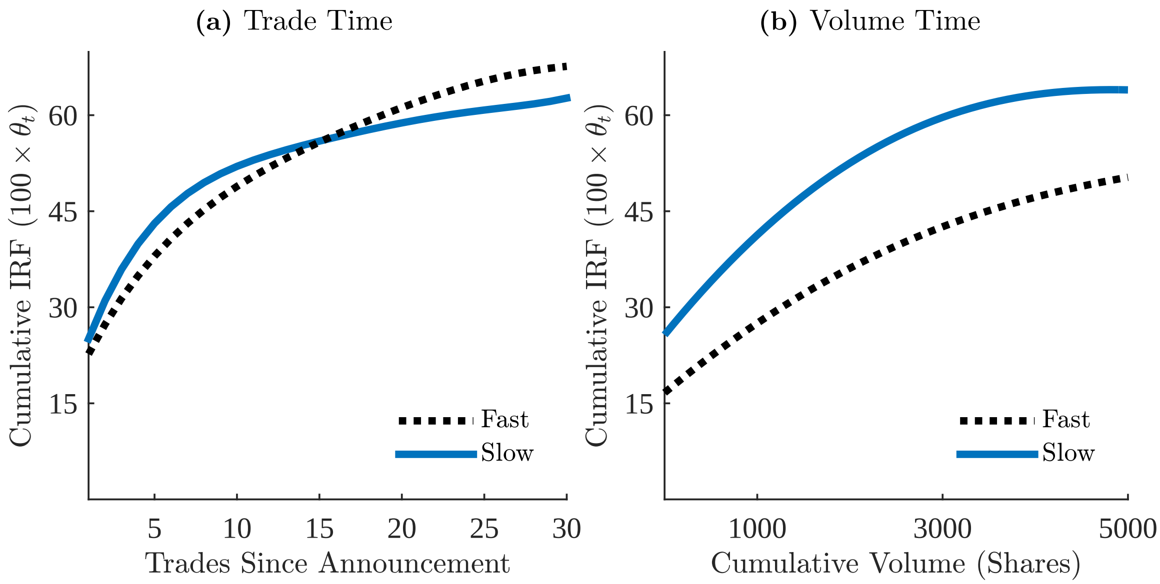

Trade 0 is the last trade prior to the earnings press release. Trade 1 is the first trade post announcement. In the left panel, firms are sorted by duration from announcement to thirtieth trade, within year, quarter, and size quintile. In the right panel, firms are sorted by the number of trades required to reach 3000 shares cumulative volume.

domain to $t \in [100,\ 10000]$ and use a polynomial approximation.[21] Allowing Fast and Slow events to have a different impulse response leads to the following specification:

$$\begin{aligned} \theta(t, \text{Slow}) = & \left(a_1 + b_1 t + c_1 t^2 \cdots\right) \times [\text{Slow} = 1] \\ & + \left(a_0 + b_0 t + c_0 t^2 \cdots\right), \end{aligned} \tag{6}$$

where the maximum degree is to be determined. In practice I use a fourth order polynomial but the results are similar using third or fifth order. As instruments I include $\left[\begin{array}{cccc} 1 & t & t^2 & \cdots \end{array}\right] \otimes \left[\begin{array}{cc} 1 & \mathbf{1}(\text{Slow} = 1) \end{array}\right] s_i$. Estimation is still by two stage least squares, but now observations at all horizons $t > 0$ are pooled together. To validate this methodology I reestimate the functions shown in Figure 6a using the continuous specification.

[21]By the Stone-Weierstrauss theorem, if $\theta(t)$ is continuous it can approximated arbitrarily well by a polynomial.

Figure 7a shows trade time impulse response functions separately for Fast and Slow firms, with speed defined based on time elapsed at the thirtieth trade. The functions look like smoothed versions of the curves in Figure 6a, confirming that the polynomial approximation works well. Figure 7b shows the volume time time impulse response functions separately for Fast and Slow firms, with speed defined based on number of trades required to reach 3000 shares. Unlike Figure 7a, the graph shows substantial heterogeneity; firms that trade slowly (fewer shares per trade) have substantially more price discovery per share traded. The difference is statistically significant at all horizons with p-value less than 5% (not shown), confirming that price discovery is not invariant when measured in volume time.

5 Robustness

5.1 Choice of Instruments

Throughout this paper I use a single instrument, $s_{i,q}$, as the exogenous instruments and do not include calendar time fixed effects when estimating θ. This contrasts with the standard approach of sorting observations into discrete bins and forming portfolios. Differencing the "high" and the "low" portfolio returns largely removes time-variation in aggregate returns, similar to using time fixed effects in a panel regression. I reestimate the θ_t using alternative instruments, with and without fixed calendar month effects:

- I include higher order terms s^2 and s^3
- As in DellaVigna and Pollet (2009), I sort earnings surprises each calendar quarter into deciles and construct a set of dummy variables for decile inclusion:[22] Deciles $= \left[\begin{array}{ccc} \mathbf{1}_{\text{Decile}=1} & \cdots & \mathbf{1}_{\text{Decile}=10} \end{array} \right]'$
- To avoid potential overfitting in the first stage, I construct the simple binary instrument High $= \mathbf{1}_{\text{Decile}\geq 6}$ which is simply an indicator for the top half of the surprise distribution (each quarter). Importantly, θ estimated using such a

[22] Results are unchanged using decile breakpoints calculated from announcements in the previous calendar quarter.

binary instrument is equivalent to the intuitive ratio Equation 1 and does not assume the linearity imposed in Equation 2

Figure 8 shows the tick-by-tick impulse response function using the three sets of instruments. The estimated functions are quantitatively similar to the estimates in Figure 3 and yield the same qualitative conclusions.

Figure 8: θ_t Using Alternative Instruments

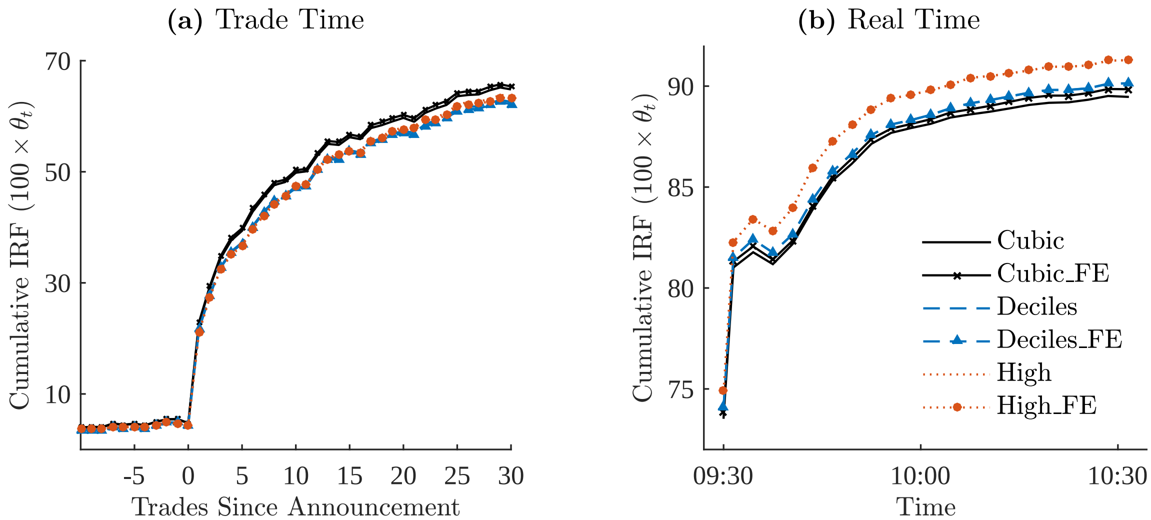

5.2 Linearity

Equation 2 postulates a perfectly linear relationship between $E[R_t | \Delta V]$ and ΔV. Some authors claim that arbitrageurs face short-sale constraints and this should slow the initial speed of price discovery for negative value shocks (Chan, 2003). This implies θ_t should initially increase faster for positive ΔV. As a second test, I estimate Equation 2 separately for positive and negative surprises. Figure 9a shows the estimated functions; the positive and negative impulse response functions are quite similar for the first few trades. As a further test of linearity I separately estimate θ_t for small and large announcements. Small announcements are those between the 25th and 75th percentiles of the surprise distribution, the "Middle". The remaining announcements are large, the "Tails". Figure 9b shows that θ_t is nearly identical for

Figure 9: Linearity

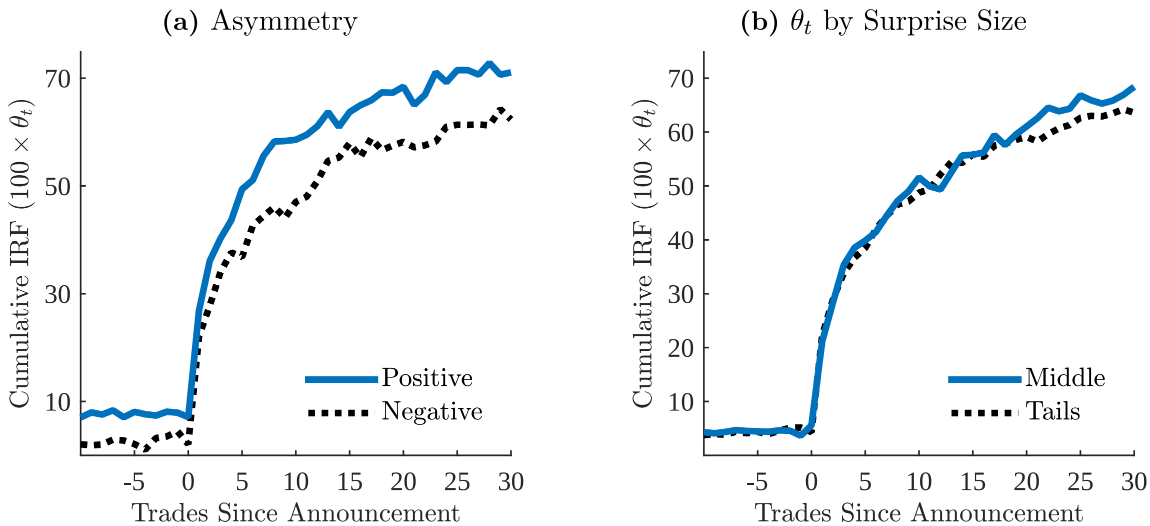

Trade 0 is the last trade prior to the earnings release. Trade 1 is the first trade post announcement. In the left panel, θ_t is estimated separately for positive and negative earnings surprises. In the right panel, θ_t is estimated separately for small and large earnings surprises.

both small and large announcements, supporting the linear specification in Equation 2. Overall I do not find compelling evidence of asymmetry in the impulses response, though these methods may lack sufficient power.

5.3 Before Market Open

The estimates in the main body of the paper use only announcements which occur after the market close (AMC). Below I also estimate the impulse response functions for announcements which occur before the regular market open (BMO). I use $T = C4$ for these announcements based on analysis similar to that in Table 1. Figure 10 shows the first day impulse response for both BMO and AMC announcements. Remarkably they are nearly identical for the first thirty minutes of the regular trading day, but then diverge.[23] Still, the general patterns are qualitatively similar, including the non-monotonicity observed near the close of trading. Table 7 decomposes price discovery for BMO announcements from C1 to C4 and shows that after the first day, most of the remaining price discovery occurs during regular trading hours.

[23]The differences are statistically significant.

Figure 10: Clock Time Cumulative IRF

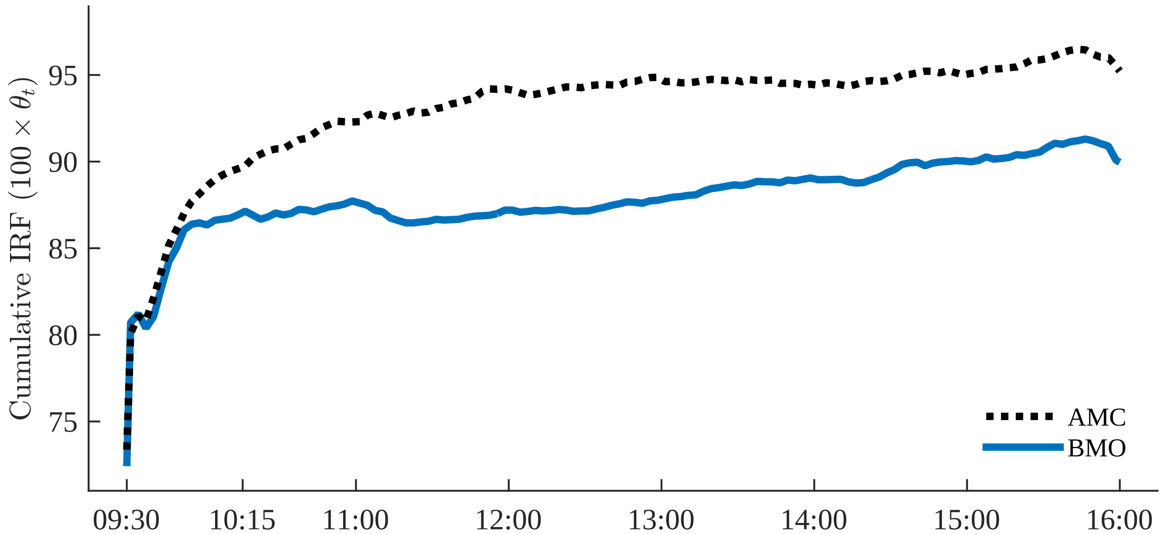

6 Conclusion

Traditional measures of the path of price discovery are based either on measures of excess volatility or return predictability. In this paper I use an alternative dimensionless measure motivated by theoretical models of how information is incorporated into prices. Using high-frequency data, I find substantial cross-sectional variation in the speed of price of price discovery when measured in "clock time." This heterogeneity, however, essentially disappears when measured in "trade time." Consistent with previous work work, I find that trade time is well approximated by the accumulation of transactions, not share volume, dollar volume, or turnover. Importantly, the *invariance* I document in trade time crucially depends on the natural rescaling inherent in my measure of price discovery. These findings challenge the traditional view that trading is not important for price discovery in response to public information, but support the hypothesis of Kyle and Obizhaeva (2016b) that "information flows take place in the same business time as trading."

Table 7: Regular vs Extended Trading

Each column displays the change in cumulative impulse response (in percent) measured over the given horizon. For example, the third column reports 7.87 for the period 10:00-3:45, meaning the cumulative impulse response increases by 7.87% during this time frame.

Standard errors are adjusted for heteroskedasticity and are clustered by firm and day of announcement. Absolute t statistics in parentheses.

	C1→O2	O2→C2	C2→O3	O3→C3	C3→O4	O4→C4
$\theta_{\omega:t}$	2.31**	7.21**	0.57	4.12**	0.01	0.88
	(3.51)	(6.15)	(1.09)	(3.49)	(0.01)	(0.79)
N	15374	15374	15374	15374	15374	15374
# Firms	1643	1643	1643	1643	1643	1643
# Year× Quarter	24	24	24	24	24	24

* $p < 0.05$, ** $p < 0.01$

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

Ané, T. and H. Geman (2000). Order flow, transaction clock, and normality of asset returns. *The Journal of Finance 55*(5), 2259–2284.

Back, K. and H. Pedersen (1998). Long-lived information and intraday patterns. *Journal of Financial Markets 1*(3), 385–402.

Barclay, M. J. and T. Hendershott (2003). Price Discovery and Trading After Hours. *Review of Financial Studies 16*(4), 1041–1073.

Barclay, M. J. and T. Hendershott (2004). Liquidity Externalities and Adverse Selection: Evidence from Trading after Hours. *The Journal of Finance 59*(2), 681–710.

Barclay, M. J. and T. Hendershott (2008). A comparison of trading and non-trading mechanisms for price discovery. *Journal of Empirical Finance 15*(5), 839–849.

Beechey, M. J. and J. H. Wright (2009). The high-frequency impact of news on long-term yields and forward rates: Is it real? *Journal of Monetary Economics 56*(4), 535–544.

Bernard, V. and J. Thomas (1990). Evidence that Stock Prices do not Fully Reflect the Implications of Current Earnings for Future Earnings. *Journal of Accounting and Economics 13*(4), 305–340.

Bessembinder, H. (1999). Trade Execution Costs on Nasdaq and the NYSE: A Post-Reform Comparison. *Journal of Financial and Quantitative Analysis 34*(3).

Brandt, M. W. and K. A. Kavajecz (2004). Price Discovery in the US Treasury Market: The Impact of Orderflow and Liquidity on the Yield Curve. *The Journal of Finance 59*(6), 2623–2654.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2012). Robust Inference with Multiway Clustering. *Journal of Business & Economic Statistics*.

Chan, W. (2003). Stock price reaction to news and no-news: Drift and reversal after headlines. *Journal of Financial Economics 70*(2), 223–260.

Chordia, T., A. Goyal, G. Sadka, R. Sadka, and L. Shivakumar (2009). Liquidity and the Post-Earnings-Announcement Drift. *Financial Analysts Journal*, 18–32.

Clark, P. K. (1973). A subordinated stochastic process model with finite variance for speculative prices. *Econometrica: journal of the Econometric Society*, 135–155.

Copeland, T. E. and D. Galai (1983). Information Effects on the Bid-Ask Spread. *the Journal of Finance 38*(5), 1457–1469.

Cushing, D. and A. Madhavan (2000). Stock Returns and Trading at the Close. *Journal of Financial Markets 3*(1), 45–67.

DellaVigna, S. and J. Pollet (2009). Investor Inattention and Friday Earnings Announcements. *The Journal of Finance 64*(2), 709–749.

Easley, D., N. M. Kiefer, and M. O'Hara (1996). Cream-Skimming or Profit-Sharing? The Curious Role of Purchased Order Flow. *The Journal of Finance 51*(3), 811–833.

Easley, D., N. M. Kiefer, and M. O'Hara (1997a). One Day in the Life of a Very Common Stock. *Review of Financial Studies 10*(3), 805–835.

Easley, D., N. M. Kiefer, and M. O'Hara (1997b). The Information Content of the Trading Process. *Journal of Empirical Finance 4*(2), 159–186.

Evans, M. and R. K. Lyons (2002). Order Flow and Exchange Rate Dynamics. *Journal of Political Economy 110*(1), 170–180.

Evans, M. D. and R. K. Lyons (2008). How is macro news transmitted to exchange rates? *Journal of Financial Economics 88*(1), 26–50.

Fama, E. (1970). Efficient Capital markets: A Review of Theory and Empirical Work. *The Journal of Finance 25*(2), 383–417.

Fleming, M. J. and E. M. Remolona (1999). Price Formation and Liquidity in the US Treasury Market: The response to Public Information. *The Journal of Finance 54*(5), 1901–1915.

Gerlach, J. and J. Lee (2011). Is the Opening Price Efficient? Evidence from After-Hours Earnings Announcements.

Glosten, L. R. and P. R. Milgrom (1985). Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders. *Journal of Financial Economics 14*(1), 71–100.

Green, J. R. . (1973). Information, Efficiency, and Equilibrium. Harvard Institute of Economic Research Discussion Paper No. 284.

Green, T. C. (2004). Economic News and the Impact of Trading on Bond Prices. *The Journal of Finance 59*(3), 1201–1234.

Hayek, F. (1945). The Use of Knowledge in Society. *The American Economic Review 35*(4), 519–530.

Hong, H. and J. Stein (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of Finance 54*(6), 2143–2184.

Hou, K. (2007). Industry Information Diffusion and the Lead-lag Effect in Stock Returns. *Review of Financial Studies 20*(4), 1113–1138.

Huang, R. D. and H. R. Stoll (1996). Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE. *Journal of Financial Economics 41*(3), 313–357.

Jones, C. M., G. Kaul, and M. L. Lipson (1994). Transactions, Volume, and Volatility. *Review of Financial Studies 7*(4), 631–651.

Kelley, E. K. and P. C. Tetlock (2013, May). Why Do Investors Trade?

Kothari, S. (2001). Capital markets research in accounting. *Journal of Accounting and Economics 31*(1), 105–231.

Kyle, A. and A. Obizhaeva (2016a). Dimensional analysis and market microstructure invariance.

Kyle, A. S. (1985). Continous Auctions and Insider Trading. *Econometrica 53*, 1315–1335.

Kyle, A. S. and A. A. Obizhaeva (2016b). Market microstructure invariance: Empirical hypotheses. *Econometrica 84*(4), 1345–1404.

Mandelbrot, B. and H. M. Taylor (1967). On the distribution of stock price differences. *Operations research 15*(6), 1057–1062.

Roll, R. (1984). A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market. *The Journal of Finance 39*(4), 1127–1139.

Thompson, S. B. (2011). Simple formulas for standard errors that cluster by both firm and time. *Journal of Financial Economics 99*(1), 1–10.

Vega, C. (2006). Stock price reaction to public and private information. *Journal of Financial Economics 82*(1), 103–133.

Appendix

A Identification

Let k be the dimension of C. Let $e_{it} = R_{it} - \mathbb{E}\left[R_{it} | \mathcal{I}_{i0},\ \text{News}_i\right]$ be the forecast error at time t conditional on news. By iterated expectations, any $z_i \in \{\mathcal{I}_{i0},\ \text{News}_i\}$ is uncorrelated with $e_{i,t}$.

$$R_{it} = \mathbb{E}\left[R_{it} | \mathcal{I}_{i0}\right] + \left(\delta_t' C_i + \eta_{it}\right) \Delta V_i + e_{it} \tag{7}$$

$$R_{iT} = \mathbb{E}\left[R_{iT} | \mathcal{I}_{i0}\right] + \Delta V_i + e_{iT} \tag{8}$$

Substituting the second into the first we have

$$R_{it} = \mathbb{E}\left[R_{it} | \mathcal{I}_{i0}\right] + \left(\delta_t' C_i + \eta_{it}\right) \left[R_{iT} - \mathbb{E}\left[R_{iT} | \mathcal{I}_{i0}\right] - e_{iT}\right] + e_{it} \tag{9}$$

$$= \delta_t' \left(C_i R_{iT}\right) + \left(e_{it} - \delta_t' C_i e_{iT}\right) + \left(\mathbb{E}\left[R_{it} | \mathcal{I}_{i0}\right] - \left(\delta_t' C_i\right) \mathbb{E}\left[R_{iT} | \mathcal{I}_{i0}\right]\right) + \eta_{it} \Delta V_i \tag{10}$$

A regression of R_{it} on R_{iT} clearly suffers from endogeneity since R_{iT} is likely correlated with the second and third "error" terms. Recall that $\mathbb{E}\left[e_{it} | \mathcal{I}_{i0},\ \text{News}_i\right] = 0$ and hence any instrument $z_i \in \{\mathcal{I}_{i0},\ \text{News}_i\}$ is uncorrelated with the second term. Identification requires that after partialing out some controls X (which span a constant) the instruments are uncorrelated with $\mathbb{E}\left[R_{it} | \mathcal{I}_{i0}\right]$ and $\left(\delta_t' C_i\right) \mathbb{E}\left[R_{iT} | \mathcal{I}_{i0}\right]$. Since $X \in \mathcal{I}_0$, it is orthogonal to ΔV and e. Further, the instruments must be uncorrelated with $\eta_{it} \Delta V_i$ (the product of unobserved heterogeneity in speed and the value shock).

First, in practice any potential covariance is likely to unimportant since the drift terms themselves are likely to be negligible over horizons considered. In what follows, assume that the unconditional expectation (across events) of $\mathbb{E}\left[R_{it} | \mathcal{I}_{i0}\right]$ is zero for all t. This is wlog since we include a constant in X. Further, covariances are *partial* covariances (across events) after orthogonalizing with respect to X. Given the endogenous regressors $C_i R_{iT}$, we require at least K instruments, Z_i, which satisfy for all $k \in \{1 \ldots K\}$: (1) $\text{cov}\left(Z_{ik},\ \mathbb{E}\left[R_{it} | \mathcal{I}_{i0}\right]\right) = 0$, (2) $\text{cov}\left(Z_{ik},\ \left(\delta_t' C_i\right) \mathbb{E}\left[R_{iT} | \mathcal{I}_{i0}\right]\right) = 0$, and (3) $\text{cov}\left(Z_{ik},\ \eta_{it} \Delta V_i\right) = 0$.

First consider conditions (1) and (2). In my estimation, the instruments take the form $Z_i = s_i C_i$. Notice that

$$\text{cov}\left(Z_{ik},\ \left(\delta_t' C_i\right) \mathbb{E}\left[R_{iT} | \mathcal{I}_{i0}\right]\right) = \mathbb{E}\left(s_i C_{ik} C_i' \delta_t \mathbb{E}\left[R_{iT} | \mathcal{I}_{i0}\right]\right)$$

$$= \text{cov}\left(s_i C_{ik} C_i' \delta_t,\, \mathbb{E}\left[R_{iT} | \mathcal{I}_{i0}\right]\right)$$
$$= \text{cov}\left(s_i C_{ik} C_i',\, \mathbb{E}\left[R_{iT} | \mathcal{I}_{i0}\right]\right) \delta_t$$

The conditions become $\forall t, k, m$ $\text{cov}\left(s_i C_{ik},\, \mathbb{E}\left[R_{it} | \mathcal{I}_{i0}\right]\right) = 0$ and $\text{cov}\left(s_i C_{ik} C_{im},\, \mathbb{E}\left[R_{it} | \mathcal{I}_{i0}\right]\right) = 0$. Essentially, products of the earnings surprise and up to two cross-sectional variables (with replacement) must be uncorrelated with prior drift at all horizons. Further, notice that if the C are a set of dummy variables, the second condition is redundant so identification requires only $\text{cov}\left(s_i C_{ik},\, \mathbb{E}\left[R_{it} | \mathcal{I}_{i0}\right]\right) = 0$. Since the C_{ik} are "group membership" dummies, the condition can be rephrased as: conditional on group identity $k \in \{1 \ldots K\}$, the earnings surprise is uncorrelated with prior drift, $\text{cov}\left(s_i,\, \mathbb{E}\left[R_{it} | \mathcal{I}_{i0}\right]\right) = 0$. Implicitly, this is the identification assumption used in any event study, that the observed "news" is truly news.

Finally, the condition $\text{cov}\left(Z_{ik},\, \eta_{it} \Delta V_i\right) = 0$ typically arises in settings with unobserved heterogeneity in "treatment effects." A sufficient condition for this to be true is independence of η_{it} and ΔV_i, or independence of the speed of price discovery and the size of the shock, conditional on the observed instruments. Note that since the instruments themselves may be correlated with ΔV, this restriction does not rule out all dependence between size and speed. It does, however, rule dependence between size and *unobserved* speed.