The Value of Information: A Puzzle

Ohad Kadan — *Arizona State University*

Asaf Manela — *Washington University in St. Louis*

January 2026

Abstract

We show that under mild assumptions, the total value of information to informed traders in the market can be measured by the covariance between price changes and order flow. This covariance captures noise trader losses, which equal informed trader gains when market making is competitive. We estimate the value of information using high frequency data on US equities at about $3.5 million per year for the average stock. The aggregate value of information is about 0.04% of market cap, which is considerably lower than the 0.67% in fees investors pay each year searching for superior returns (French 2008). We discuss potential resolutions for these puzzling findings.

JEL classification: G10, G12, G14, D80, D82

Keywords: value of information, active asset management fees, liquidity, order flow

We thank Julie Fu, Nicolae Gârleanu, Péter Kondor, Taisiya Sikorskaya, seminar participants at LBS, LSE and Washington University in St. Louis, and participants at the 2026 ASU-HKU Conference for helpful comments. Claude Opus 4.5 and GPT-5.2 provided research assistance. E-mail: ohad.kadan@asu.edu or amanela@wustl.edu.

1 Introduction

How much would you pay for information about tomorrow's closing stock price? What may at first sound like an arbitrage opportunity with infinite value, is in fact bounded in most markets due to price impact. Informed investors with foresight of a future stock price would act strategically and gradually trade to shade their superior information between the orders of other uninformed traders (Kyle 1985). The expected gain of the informed, which we call the value of information, provides an upper bound on the value added by active asset managers who promise, for a fee, to beat the market using their superior information. In what follows, we estimate the aggregate value of information for all stocks traded in the US and for all potentially informed investors combined. We find that it is surprisingly small compared to the aggregate amount of fees paid by investors to active asset managers (e.g., French 2008).

Our approach to estimating the value of information relies on mild assumptions, ones satisfied by a broad class of theoretical models on strategic informed trading. These models include the Kyle (1985) and Back (1992) models of monopolistic informed traders, the Back, Cao, and Willard (2000) generalization to multiple informed investors with imperfect signals, and the Collin-Dufresne and Fos (2016) extension to stochastic noise trading. In such models, informed investors gain at the expense of noise or liquidity traders who lose, while competitive market makers break even on average. Noise trader losses on average equal informed trader gains when market making is competitive, and provide an upper bound on the value of information when market makers have market power (Treynor 1971). Rather than a hypothetical signal, mean noise trader losses measure the value of actual information in the market using the revealed preferences of market makers.

While measuring the average gains of the informed is often complicated, measuring the average losses of noise traders is surprisingly tractable. Noise traders lose every time they submit a market buy order dZ and the price increases dP due to price impact (or a

bid-ask spread) set by market makers. They also lose if they sell and price decreases. It has long been known, therefore, that the total value of information, Ω, equals noise traders' expected losses over the trading period T,

$$\Omega = E \int_0^T dP dZ, \tag{1}$$

which is simply the covariance (formally the expected quadratic covariation) between price changes and noise trades. Price changes dP are clearly observable. However, the literature has so far struggled with the measurement of noise trading dZ. If market makers cannot distinguish between informed trades and noise trades, how could an econometrician possibly do so?

Our key insight is that observing noise trading is unnecessary for measuring the covariance in Equation (1). The reason is that in this class of models, the informed investor's trading strategy is locally orthogonal to price changes. Informed traders try to disguise their trades to avoid a strong price impact by smoothing their trades over time. As a result, informed trading dX is "smooth," while noise trading dZ is "jittery." Consequently, the covariance between price changes and changes in noise trading exactly equals the covariance between price changes and order flow $dY = dX + dZ$, as long as it is measured at a high enough frequency. This, in turn, implies that the total value of information (1) takes the form

$$\Omega = E \int_0^T dP dY. \tag{2}$$

Thus, measuring order flow dY, which is observable, suffices for quantifying the aggregate losses of uninformed noise traders, and therefore to measure the value of information.

Armed with this simple expression, we estimate the value of information daily with a discrete version of Equation (2), using intraday data for US publicly-listed common stocks over the September 2003 to December 2024 period. The idea is to measure an investor's

willingness to pay for information on the average stock on the average trading day. We find that the average daily value of information is $13,700, which annualizes to an average of $3.5 million for becoming informed on an average stock each day.

The value of information we estimate is surprisingly small. The aggregate annualized value of information for all US stocks combined is only 0.04% of market capitalization. This value is the most active investors and asset managers can hope to gain by seeking superior information to forecast stock prices. For comparison, French (2008) estimates that investors pay active asset managers about 0.67% of market capitalization in fees each year searching for superior returns. As Gârleanu and Pedersen (2018) shows, in a general equilibrium for both assets and asset management, the value of information equals management fees plus the cost of searching for good managers. We find it puzzling therefore that investors would pay so much more in fees than active asset management can possibly generate.

We initially met this finding with skepticism. We therefore devote considerable space to validation and robustness tests. In the timeseries, we find that the value of information rises during turbulent times consistent with Kadan and Manela (2025). In the cross-section, we find that the value of information is higher for large-growth stocks, consistent with Farboodi et al. (2021), and higher for momentum stocks, consistent with Kadan and Manela (2025). On earnings announcement days, when we expect informed investors to have some of the most valuable information, we find that the value of information rises more than 3-fold. We further find that the magnitude of the value of information hardly changes when we vary the frequency of observation, the way we sign trades as buys or sells, and if we include penny stocks.

In the last part of the paper we discuss potential resolutions to the value of information puzzle. We consider trends in the costs of active investing and other non-informational services provided by asset managers. We consider strategies concentrating information acquisition on days with high value of information. We discuss the role of risk aversion,

the role of partial information, and the role of multiple informed investors. We discuss the role of non-competitive market makers and non-competitive asset managers. We generalize our results to discrete time, allowing for jumps, and relax the assumption that price changes are orthogonal to informed trading. Finally, we discuss the role of behavioral investors.

Our paper is the first to estimate the value of information indirectly, by measuring the losses of noise traders, and the first to point out the conditions under which the covariation (i.e., quadratic covariation) between price changes and order flow measures exactly this quantity. The generality and simplicity of our measure allow us to reliably quantify the value of information and to document its puzzling magnitude relative to active asset management fees. For this purpose, it is paramount that it has easy to interpret units.[1] The groundwork for our analysis lies in the classic contributions of Grossman and Stiglitz (1980) and Kyle (1985), who evaluate the value of information within a rational-expectations framework. Epstein, Farhi, and Strzalecki (2014) and Kadan and Manela (2019) examine the value of macroeconomic information under Epstein and Zin (1989) preferences, which lets them distinguish between the instrumental and psychic components of information. Croce, Marchuk, and Schlag (2023) quantifies the relevance of heterogeneity in the timing of exposure of cash flows to aggregate shocks for returns. Ai et al. (2026) extend the option-based methodology of Kadan and Manela (2019) to accommodate generalized risk preferences. Farboodi et al. (2021) and Farboodi et al. (2025) compute the value of information implied by a structural noisy rational-expectations model and show it is higher for large growth firms. Bogousslavsky, Fos, and Muravyev (2024) develop a machine learning-based measure of informed trading intensity that captures the prevalence and timing of informed trading.[2] Kadan and Manela (2025) show that the ratio

1. See Kyle and Obizhaeva (2016), Kyle and Obizhaeva (2018), Kyle, Obizhaeva, and Wang (2018), and Kyle and Obizhaeva (2019) for careful work on microstructure invariants and units of measurement.

2. The CARA–normal (mean–variance) noisy rational-expectations framework underpins a large literature on information choice; see Veldkamp (2011) for a survey and Malamud and Rostek (2017) and Dávila

of fundamental variance to price impact provides an approximation for the value of information in the Back, Cao, and Willard (2000) model with multiple differentially informed investors, and estimate it using high-frequency stocks data. Here we provide a precise expression (rather than an approximation) for the value of information in a more general setting which incorporates this previous work as a special case, and use it to document its puzzling magnitude relative to active asset management fees.

Our paper also relates to research on the informational efficiency of financial markets (Fama 1970). Recent work documents that price informativeness has risen since the 1960s, especially among large growth firms (Bai, Philippon, and Savov 2016; Farboodi et al. 2021). Dávila and Parlatore (2025) estimate informativeness by regressing prices on fundamentals and find it is greater for large stocks with high turnover, idiosyncratic volatility, institutional ownership, and analyst coverage. Kacperczyk, Sundaresan, and Wang (2021) report that higher foreign ownership boosts price informativeness, and Dávila and Parlatore (2023) examine its relation to volatility. We contribute a simple measure of the value of information that incentivizes market participants to acquire private signals and, through trading, enhance capital allocation and managerial decisions (Bond, Edmans, and Goldstein 2012; Brogaard, Ringgenberg, and Sovich 2019; Goldstein 2023). An interpretation of our finding that the value of information is modest, is that the equity market is highly efficient.

The paper proceeds as follows. Section 2 provides the theoretical foundation for our measure of the value of information. Section 3 discusses its estimation. Section 4 discusses the value of information puzzle and its potential explanations. Section 5 concludes. Robustness tests are in the Appendix.

and Parlatore (2021) for recent contributions. However, empirical implementation is challenging because stock prices are neither stationary nor Gaussian, and recent studies suggest that departures from CARA utility can be consequential (Savov 2014; Breon-Drish 2015; Malamud 2015). Our estimate by contrast is agnostic to the functional form of the utility function.

2 Theory

In this section, we use a generic version of continuous-time Kyle-Back models of trade with asymmetric information to develop a general closed-form expression for the value information. The main novel point we are making is that under standard assumptions, the exact functional form of the equilibrium is not needed to calculate the value of information, which takes a simple close form that is easy to estimate.

2.1 Setup

An asset is traded by three types of agents: informed traders, uninformed noise traders, and market makers. Informed and uninformed submit their buy and sell orders, and the market makers set prices to clear the market.

The asset has a random fundamental (terminal) value v with a continuous distribution function F. Similar to Back (1992), we do not make specific functional assumptions on F, which can accommodate a variety of different functional forms. The asset is traded continuously over a time interval $[0, T]$. One or more informed traders get some information about v at time $t = 0$. This information may be perfect as in Kyle (1985) or imperfect as in Back, Cao, and Willard (2000). Their cumulative orders at time t are denoted by X_t, assumed to be a semi-martingale.

Noise traders demand Z_t shares on net through time t. The cumulative noise trading process is given by

$$dZ_t = \sigma_z(t) dB_t, \tag{3}$$

where B_t is a Wiener process and $\sigma_z(t)$ is the instantaneous volatility of noise trading, which may be stochastic as in Collin-Dufresne and Fos (2016). Market makers cannot tell apart the informed and uninformed trades but observe the total order flow $Y_t = X_t + Z_t$. We denote the instantaneous variance of order flow by σ_y^2. Market makers set the price P_t

based on Y_t, such that the price process is a semi-martingale.[3]

The noise traders' losses stem from prices moving against them when they trade. For example, if they submit an order to buy $dZ_t > 0$ shares and the price per share increases by $dP_t > 0$, then they effectively pay a "slippage" cost of $dP_t dZ_t > 0$. Given this, the total expected loss of noise traders is $E \int_0^T dP_t dZt$.

Our first key assumption is that market makers are competitive and thus earn zero profits on average. As a result, the value of information, which is the expected profit to informed traders (denoted by Ω), must equal the expected loss to noise traders.

Assumption 1 (Competitive Market Making)**.** *The total value of information to informed traders equals the expected loss to noise traders,*

$$\Omega = E \int_0^T dP_t dZ_t.$$

Assumption 1 is standard in Kyle-Back models such as Kyle (1985), Back (1992), Back, Cao, and Willard (2000), and Collin-Dufresne and Fos (2016). In such models, the value of information may be associated with model specifics such as the number of informed traders, the accuracy and correlation of their signals, and the distribution of the fundamental value, all of which are hard to measure empirically and turn out to be not necessary for our purposes. Assumption 1 says that in order to estimate the value of information what we need is to estimate the losses to noise traders, which, in turn, already capture all the above model parameters. Our fundamental insight is that, empirically, measuring the expected losses of noise traders is considerably simpler than measuring the expected profits of informed traders directly.

Our second assumption states that the informed trading process is stochastically uncorrelated with the price process.

3. The technical assumption that X_t and P_t are semi-martingales is standard in Kyle-Back models. It is required to ensure that the stochastic integrals below are well-defined. Note that Z_t is also a semi-martingale by virtue of the Wiener process B_t.

Assumption 2 (Orthogonal Informed Trading)**.** *Informed traders' strategy is orthogonal to the price process, i.e.,* $\int_0^T dX_t dP_t = 0$.

Formally, this means that the quadratic covariation of X_t and P_t is zero. Assumption 2 is a feature of all continuous-time Kyle-Back style models we are aware of including Kyle (1985), Back (1992), Back, Cao, and Willard (2000), and Collin-Dufresne and Fos (2016). In these models, informed trading is locally deterministic and therefore has zero quadratic covariation with the price process.

The intuition behind Assumption 2 is compelling and is now standard in Kyle-Back type models. Economically, the quadratic covariation, $\int_0^T dX_t dP_t$, captures the cumulative losses to the informed traders from the price impact of their trades. This magnitude tends to be positive by the fact that market makers try to tease out the informed trades from noise trades through observed order flow. Back (1992) shows that informed traders optimally choose a smooth trading strategy that minimizes the cumulative price impact to their trades, thereby setting $\int_0^T dX_t dP_t$ to zero, as stated in Assumption 2. Thus, this assumption reflects that informed traders are trying to disguise their trades to avoid a strong price impact by smoothing their trades over time.

As Back (1992) notes, the condition $\int_0^T dX_t dP_t = 0$ essentially means that informed trades during small time increments of Δt are of the order Δt. By contrast, noise trading is driven by a Brownian motion, and therefore the magnitude of noise-trading shocks during small time increments of Δt is of order $\sqrt{\Delta t}$. Thus, informed trading is "smooth," while noise trading is "jittery." Variations in prices are then dominated by the jittery nature of noise trading, and the quadratic covariation of prices with informed trading vanishes. If informed traders were to add randomness to their trades in the form of a Brownian motion with increments of order $\sqrt{\Delta t}$ (like noise traders), their orders would contribute to the diffusive part of prices. But such behavior is suboptimal in the Kyle-Back environment because it would render $\int_0^T dX_t dP_t$ being positive. The informed traders' optimal strategy

in those models is therefore to spread trades smoothly over time, which is exactly what forces the quadratic covariation of X_t with P_t to vanish.

2.2 The value of information

As noted, Assumptions 1 and 2 are satisfied by essentially all continuous-time Kyle-Back type models of trading under asymmetric information. While these models may differ in terms of the number of informed traders and the nature of private information, these two assumptions summarize standard and well-accepted features of these models.

Note that the stochastic loss by noise traders $\int_0^T dP_t dZ_t$ can be written as:

$$\int_0^T dP_t dZ_t = \int_0^T dP_t(dY_t - dX_t) = \int_0^T dP_t dY_t, \tag{4}$$

where the last equality is a consequence of Assumption 2. It follows that the total expected loss to noise traders is given by the $E\int_0^T dP_t dY_t$, which by Assumption 1 is the value of information. We have established the following simple but powerful result.

Proposition 1 (Value of information)**.** *Under Assumptions 1 and 2, the value of information over an interval of length T is given by*

$$\Omega = E\int_0^T dP_t dY_t.$$

The key in Proposition 1 is that we can replace noise trading changes dZ_t with order flow dY_t in the calculation of noise traders losses. While noise trading is not observable and cannot be easily inferred from data, order flow is observable and available at high frequency from standard databases, which facilitates our empirical analysis. It is important to note that noise trading and order flow are not equal to each other. The reason the substitution of dZ_t with dY_t in the proposition is valid is that under Assumption 2, both have

the same quadratic covariation with price changes. From an empirical perspective, Proposition 1 allows us to estimate the value of information by simply estimating the integral $\int_0^T dP_t dY_t$ using high frequency data.

2.3 A useful decomposition

We next show that with two additional assumptions, the value of information can be decomposed into two intuitive components. It is important to emphasize that these additional assumptions are not necessary for our estimation of the value of information. We introduce them only to facilitate interpretation.

The first additional assumption imposes additional restrictions on the trading strategies used by informed and noise traders.

Assumption 3 (Diffusive trading strategies)**.** *The informed trading strategy dX_t satisfies*

$$dX_t = \theta_t dt,$$

where θ_t is absolutely continuous. That is, dX_t has no stochastic term. Moreover, the noise trading process is given by

$$dZ_t = \sigma_z dB_t.$$

where B_t is a Wiener process and σ_z is a volatility parameter.

This assumption is common in the literature. Indeed, in all Kyle-Back models we are aware of informed trading is locally deterministic. As for the restriction on noise trading, this assumption rules out the setting of Collin-Dufresne and Fos (2016) where σ_z is itself a stochastic process. A direct implication of this assumption is that the quadratic variation of order flow is constant and equals to the quadratic variation of noise trading, $(dY_t)^2 = \sigma_z^2 dt = \sigma_y^2 dt$.

The second additional assumption is on the price process. In the classic Kyle-Back continuous-time models, prices are linear in order flow with Kyle's-lambda being the relevant slope. Critically, in those models, asset-specific order flow is the only reason for price changes. This strong assumption would suffice for our purposes. But, it is quite plausible, in fact, that other factors, orthogonal to the asset's order flow, may affect prices. For example, if a systematic risk factor, such as the market, changes in value, the price of individual stocks may change as well. To allow for that, let W_t be a Wiener process uncorrelated of B_t. The idea is that W_t represents publicly observable price-relevant information that is not captured in the order flow. We assume that prices respond to both order flow, Y_t, and to W_t in a linear fashion:

Assumption 4 (Linear prices)**.** *Order flow Y_t affects price P_t according to*

$$dP_t = \lambda_t dY_t + \sigma_w dW_t, \tag{5}$$

for some price impact function λ_t, coefficient σ_w, and a Wiener process W_t which is orthogonal (i.e., has zero quadratic covariation) with X_t and B_t.

Note that in the traditional Kyle-Back type models, price changes are assumed to be perfectly correlated with order flow. That is σ_w is set to be identically zero. Therefore, our approach incorporates these settings as a special case. Because we abstract from many model specific details, our setup allows for other public information W_t to arrive in the market and affect the asset's price as in Back, Crotty, and Li (2018). This is particularly useful from an empirical perspective, because, in reality, prices are not perfectly co-linear with order flow.

The next proposition establishes that under Assumptions 1–4, the value of information can be decomposed into a particularly intuitive form.

Proposition 2 (Value of information decomposition)**.** *Under Assumptions 1–4, the value of*

information over an interval of length T is the product of expected price impact $E[\lambda_t]$ and order flow variance $\sigma_y^2 T$:

$$\Omega = E\left[\lambda_t\right]\sigma_y^2 T. \tag{6}$$

Proof. Using Assumption 4,

$$\begin{aligned}\int_0^T dP_t dY_t &= \int_0^T \lambda_t \left(dY_t\right)\left(dY_t\right) + \int_0^T \sigma_w \left(dW_t\right)\left(dY_t\right) \\ &= \int_0^T \lambda_t \left(dY_t\right)^2 + \int_0^T \sigma_w \left(dW_t\right)\left(dX_t + dZ_t\right) \\ &= \int_0^T \lambda_t \sigma_y^2 dt, \end{aligned} \tag{7}$$

where the last equality follows from the fact that W_t is orthogonal to X_t and Z_t.

By Proposition 1,

$$\begin{aligned}\Omega &= E\left[\int_0^T dP_t dY_t\right] \\ &= E\left[\int_0^T \lambda_t \sigma_y^2 dt\right] \\ &= E\left[\frac{1}{T}\int_0^T \lambda_t dt\right]\sigma_y^2 T \\ &= E\left[\lambda_t\right]\sigma_y^2 T, \end{aligned} \tag{8}$$

as stated.

□

This result is intuitive. The value of information equals the expected aggregated losses of noise traders over time. By Proposition 1, these losses are the time-aggregated product of price changes and the order flow. Proposition 2 establishes that under linear pricing, the former aggregates to expected price impact, $E\left[\lambda_t\right]$, while the latter aggregates to order

flow variance over the interval of length T, $\sigma_y^2 T$. In our empirical analysis, we rely on this decomposition to gain economic intuition about variations in the value of information.

2.4 Discussion and connection to earlier work

Our approach in this paper is different from prior work in that we abstract from many detailed assumptions on fundamentals and instead, focus on standard features of Kyle-Back models that are crucial for the estimation of the value of information. In doing so, we are gaining simplicity and generality but at the cost of losing some details. Prior work such as Back, Cao, and Willard (2000), delved very deeply into the dynamics of informed trading, yielding important insights on how informed traders spread their trades over time. Our purpose in this paper is quite different because our sole focus is on the estimation of the value of information. As it turns out, for this purpose, the detailed dynamics of informed trading are not needed because the value of information to informed traders is simply the expected losses to noise traders (Assumption 1). Thus, once one can estimate these losses, the exact nature of the dynamics of informed traders becomes unnecessary for the calculation of the value of information. As a result, one can avoid making assumptions on the value distribution, number of informed traders, the nature and precision of the information they possess, and the underlying correlation structure.

In our setting, the challenge then becomes to estimate the losses of noise traders, which by itself is not observable. However, the quadratic covariation of noise trading with price changes is equal to that of total order flow – a consequence of Assumption 2. This, in turn, allows us to use variations in total order flow—which is empirically observable—to estimate the losses of noise traders.

In Kadan and Manela (2025), we take a different approach to estimating the value of information. We derive approximate bounds on the value of information in one specific Kyle-type model – that of Back, Cao, and Willard (2000). We then show that in that setting,

the value of information is approximately equal to the ratio of fundamental variance σ_v^2 to price impact at time zero – $\lambda(0)$. The setting offered in this paper incorporates Kadan and Manela (2025) as a special case. A major improvement with the approach provided in this paper is that we no longer rely on an approximation derived from lower and upper bounds on the value of information. Instead, the expressions obtained in Propositions 1 and 2 are exact.

Another improvement relative to much of the Kyle-Back literature is that we do not require an exact linear relation between price changes and order flow (perfect correlation). Instead, we allow for price changes to be associated with other observable sources of information. This is particularly important from an empirical perspective—in our sample, for individual stocks, the mean correlation between prices changes and order flow is only about 0.3.

3 Estimation and validation

We next explain how we estimate the value of information and validate our measure.

3.1 Data and estimation

We estimate daily values of information for each stock in the NYSE TAQ database, which covers all US publicly traded stocks. Our sample includes 8,300 common stocks over the September 11, 2003 to December 31, 2024 sample period, which we could match by trading symbol and date to CRSP/Compustat. Following Amihud (2002) we drop all stocks with a previous-day closing stock price smaller than \$5 to avoid market microstructure effects.[4]

Empirically, we interpret the daily value of information for a particular stock as the dollar amount informed investors would be willing to pay at the close of the trading day

4. In the Appendix we show how our estimates change if we keep penny stocks.

for a signal about the closing price at the end of the next trading day. Our approach captures signal precision indirectly through the losses of noise traders. All else equal, a more precise signal means higher price impact and therefore higher value of information.

Our main estimator for the value of information relies on a discrete approximation of the integral in Proposition 1. Let $P_{jt\tau}$ denote the price of asset j on day t, at time $\tau \in [0, T]$, where T denotes the time in years that passes on day t. We observe equidistant observations over time intervals $h = 1$ minute, so we estimate the sum at time T based on $N + 1$ discrete observations recorded at times τ_0, τ_1, ..., $\tau_N = Nh = T$ (We vary h in robustness analysis). We annualize the value of information by setting $T = 1/252$. Let $\Delta P_{jti} = P_{jt\tau_i} - P_{jt\tau_{i-1}}$ be the change in the price of asset j on day t over the interval $i = 1, \ldots, N$. Similarly, let $\Delta Y_{jti} = Y_{jt\tau_i} - Y_{jt\tau_{i-1}}$ be the change in cumulative signed share order flow, or simply net order flow over the same interval. For each stock j on date t, we estimate a daily annualized value of information,

$$\hat{\Omega}_{jt} = \frac{1}{T} \sum_{i=1}^{N} \Delta P_{jti} \Delta Y_{jti}. \tag{9}$$

For stationarity, we measure the value of information in two main ways. The first is in millions of dollars inflation-adjusted using the CPI to December 2024 values. The second is as a percentage of aggregate stock market capitalization. The first way is sometimes easier to interpret, while the second facilitates comparison with fees estimates such as in French (2008).

Several approaches to signing TAQ trades as buys $(+1)$ or sells (-1) have been proposed by prior work. Our baseline estimates are based on the Chakrabarty et al. (2007) approach, using the algorithms developed by Holden and Jacobsen (2014). In the Appendix we show that our qualitative and quantitative conclusions are not sensitive to this choice.

We rely on the decomposition in Proposition 2 to gain intuition for the source of varia-

tion in the value of information over time and across firms. The decomposition allows us to ask if higher value is due to higher price impact or order flow variance. To get a decomposition that does not depend on an arbitrary number of shares outstanding, we focus on annualized dollar measures, and decompose the value of information as follows. Let C be a scaling constant used to convert a dollar amount to millions of real dollars. Let $\tilde{\lambda} = \frac{E[\lambda_t]}{P_0/C}$ measure the effect of a million dollar buy order on returns. Let $\tilde{\sigma}_y = \sigma_y P_0/C$ measure the annualized volatility of order flow in millions dollars. Then, the annualized log value of information in real millions of dollars can be written,

$$\log(\Omega/C) = \log\tilde{\lambda} + \log\tilde{\sigma}_y^2. \tag{10}$$

Of course, when we use this decomposition, we restrict the sample to observations with a positive price impact coefficient.

Consistent with Assumption 4, we estimate price impact λ_{jt} using a regression of 1-minute log returns on contemporaneous share order flow,

$$r_{jti} \equiv \log\left(P_{jt\tau_i}/P_{jt\tau_{i-1}}\right) = \hat{\lambda}_{jt}\Delta Y_{jti} + w_{jti}. \tag{11}$$

We note that while in theory λ is positive, it could be negative in any finite sample. In addition, $\hat{\lambda}_{jt}$ may be subject to selection bias if traders use price-dependent strategies and cancel costly orders (Obizhaeva 2011). A large literature proposes other impact measures (Holden and Jacobsen 2014). We favor our measure because it aligns closely with the model used to interpret the data, making the value-of-information units straightforward. Finally, we use the closing price of the preceding trading day, $P_{jt-1\tau_N}$ as the initial asset price P_0.

Table 1: Information values and other summary statistics

variable	Mean	Std	Min	q25	Median	q75	Max	Obs
Information value (dPdY), $M	3.47	61.59	-1.11	0.30	1.02	3.08	52.30	14,382,920
Information value ($\lambda\sigma_y^2$), $M	3.45	61.74	-1.11	0.30	1.01	3.06	52.15	14,382,920
Orderflow Volatility, $M	45.03	98.61	0.07	2.64	10.36	38.74	662.18	14,382,920
Price impact per $M	0.11	0.38	-0.06	0.00	0.01	0.05	2.92	14,382,920
Stock price	56.49	152.30	5.00	17.84	34.18	62.75	305,971.48	14,382,920
Size (lnME)	14.23	1.82	10.46	12.92	14.13	15.41	18.93	14,382,920
Book-to-market ratio	0.60	0.45	0.03	0.28	0.50	0.80	2.59	14,382,920
Momentum	0.18	0.49	-0.66	-0.11	0.10	0.34	2.50	14,382,920
Earnings	0.01	0.11	0.00	0.00	0.00	0.00	1.00	14,382,920

Notes: Daily common stocks panel from NYSE TAQ, CRSP, and Compustat, September 2003 to December 2024. The value of information is the annualized sum of the product of share price changes times signed share order flow and reported in millions of dollars. The second line provides an alternative estimator for the value of information that is annualized price impact time order flow variance, which requires an additional assumption of linearity between price changes and signed order flow. Dollar figures are CPI adjusted to December 2024 values. Order flow volatility is the square root of the annualized one-minute signed share order flow variance. Price impact is estimated by regressing one-minute log returns on contemporaneous share order flow and divided by the previous trading day's closing stock price. Size is log market equity over the previous month. Momentum is the return over the prior 2–12 months. Book-to-market is book equity as of last June divided by market equity as of last December (Fama-French conventions). Earnings indicates an earnings announcement day. All variables are 1% winsorized.

3.2 Estimated information values and other summary statistics

Table 1 reports sample moments of the value of information as well as summary statistics for the sampled firms. Our estimates refer to the annualized value of acquiring information that is realized over a single day. Hence, the reported values represent the willingness to pay for receiving daily information on a stock for an entire year. The average annualized value of information in our sample is \$3.5 million (\$13,700 per day), with substantial dispersion: the standard deviation is about \$62 million.

Alongside our main estimator (9), derived from Proposition 1, we also present estimates based on the decomposition in Proposition 2. In that case, we first estimate λ and σ_y^2 separately for each stock–day and then take their product. The two sets of estimates are virtually identical, reinforcing that the decomposition closely approximates the general value of information.

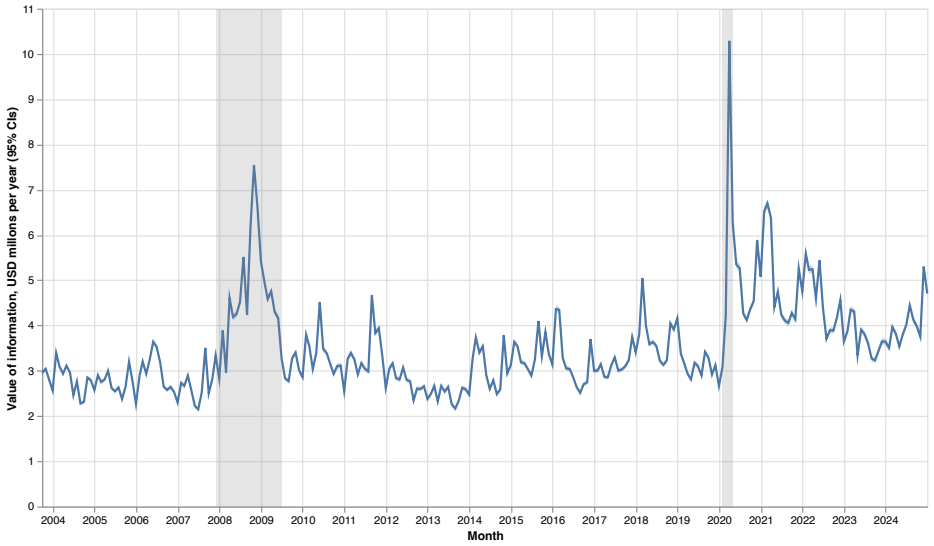
The median price impact (λ divided by the stock price) is 0.01, implying that \$1 million of order flow moves prices by 1 percent.[5] About 4 percent of stock-day observations yield a negative information value estimate. We include these when reporting information value in levels, but exclude them when reporting results in logs below. The median firm in our sample has a market capitalization of \$1 billion, a book–to–market ratio of 50 percent, and momentum indicating a 10 percent appreciation over the prior 11 months. Firms reported earnings on roughly 1 percent of sample days.

3.3 The value of information over time

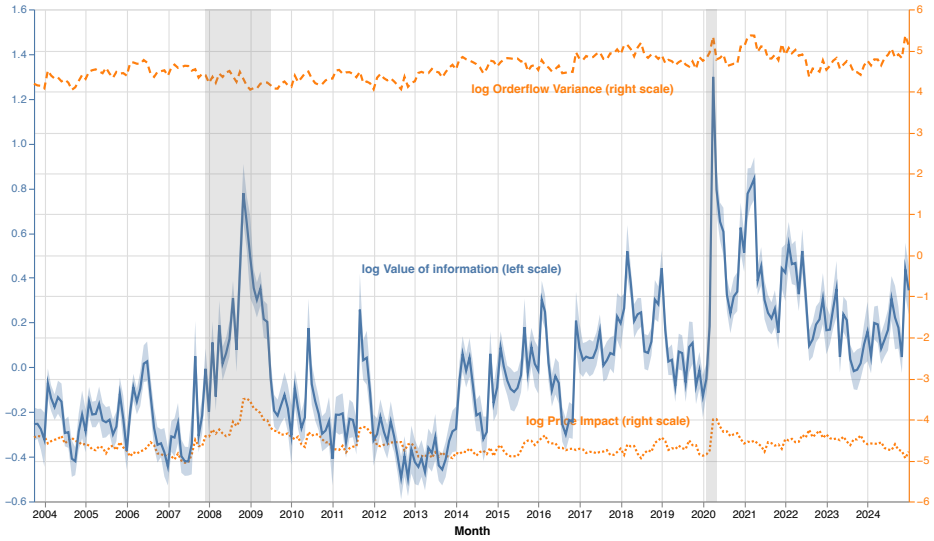
Figure 1 traces the evolution of the value of information over time. Panel (a) plots the level of $\tilde{\Omega}$, averaged across stocks and days by month. Panel (b) shows its log decomposition into log price impact and log order flow volatility, which add up to log value of informa-

5. Our median price-impact estimate aligns with Hasbrouck (2009, Table II), where a \$1 million order shifts the log price by 0.7%.

Figure 1: Value of information over time



(a) Value of information in levels



(b) Value of information in logs and its components

Notes: The solid line is the monthly information value averaged over stocks and days surrounded by its 95 percent confidence interval, based on time-clustered variance estimates. The value of information is the annualized sum of the product of share price changes times signed share order flow and reported in millions of dollars. The log information value equals log price impact plus log annualized order flow variance, so we also report the mean of log price impact (dotted line) and the mean of log annualized return variance (dashed line). Shades indicate recessions.

tion. Recessions are shaded. The figure indicates that in turbulent periods—especially recessions—order flow volatility rises and liquidity deteriorates, as seen in higher price impact. These result in higher values of information. This pattern is most pronounced in the 2007–2009 financial crisis and the 2020 Covid-19 episode, where both order flow volatility and illiquidity surge, producing a spike in the value of information.

3.4 The value of information across firms

Previous work has found that the value of information is higher for large (high market cap) and growth (low book-to-market) stocks (Farboodi et al. 2021) and for momentum (high recent returns) stocks (Kadan and Manela 2025). We confirm that this pattern shows up when using our new measure with less noise.

Table 2 studies the extent to which the value of information can be explained by commonly studied characteristics: size, book-to-market, and momentum. Panel 2a shows that large, growth and momentum stocks have higher values of information. Unlike book-to-market, which is subsumed by the inclusion of both stock and day fixed effects, large stocks and stocks that appreciated over the preceding year (high momentum stocks) have higher values of information regardless of the regression specification.

Contrasting the R-squares to the Within-R-squares, we see that these characteristics explain a substantial share of the variation in information values, even after absorbing stock and day fixed effects.

Panel 2b investigates the underlying reason for this variation using the decomposition of the value of information into price impact and order flow volatility (Proposition 2). It shows that large-momentum stocks exhibit noisier order flow and more liquid prices.

Table 2: Value of information and stock characteristics

	log Information value			
	(1)	(2)	(3)	(4)
Size	0.70***	0.70***	0.59***	0.62***
	(0.01)	(0.01)	(0.01)	(0.01)
Book-to-market	-0.29***	-0.27***	-0.07***	-0.01
	(0.02)	(0.02)	(0.01)	(0.01)
Momentum	0.29***	0.38***	0.14***	0.21***
	(0.01)	(0.01)	(0.01)	(0.01)
Date (day) FE		Yes		Yes
Stock FE			Yes	Yes
N	13,855,033	13,855,033	13,854,969	13,855,033
R^2	0.52	0.55	0.65	0.68
Within-R^2		0.54	0.10	0.10

(a) Information values are higher for large, growth and momentum stocks

	log Orderflow Variance				log Price impact			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Size	1.80***	1.82***	1.53***	1.63***	-1.11***	-1.12***	-0.95***	-1.02***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.00)	(0.00)	(0.01)	(0.01)
Book-to-market	-0.34***	-0.32***	-0.17***	-0.04**	0.06***	0.06***	0.10***	0.03***
	(0.03)	(0.03)	(0.02)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)
Momentum	0.45***	0.46***	0.29***	0.27***	-0.16***	-0.08***	-0.15***	-0.06***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.00)
Date (day) FE		Yes		Yes		Yes		Yes
Stock FE			Yes	Yes			Yes	Yes
N	13,855,033	13,855,033	13,854,969	13,855,033	13,855,033	13,855,033	13,854,969	13,855,033
R^2	0.76	0.77	0.83	0.84	0.69	0.71	0.74	0.76
Within-R^2		0.77	0.26	0.24		0.71	0.17	0.16

(b) Large-momentum stocks exhibit noisier order flow and more liquid prices

Notes: Panel (a) shows regressions of the log value of information and its components on stock characteristics. The log information value equals log price impact plus log annualized order flow variance, so in Panel (B) we regress log price impact and log annualized order flow variance on the same variables. Size is log market equity over the previous month. Momentum is the return over the prior 2–12 months. Book-to-market is book equity as of last June divided by market equity as of last December (Fama-French conventions). Observations with negative price impact are omitted. Standard errors clustered by date and stock are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.5 The value of information around earning announcement days

We next turn to studying how the value of information changes on dates in which major pieces of information are released. Intuitively, a strategic trader should be willing to pay a higher amount to learn the end-of-day stock price at the beginning of days in which information is scheduled to be released. To this end, we focus our attention on the most prominent information release dates in the life of a firm—earnings announcement dates.

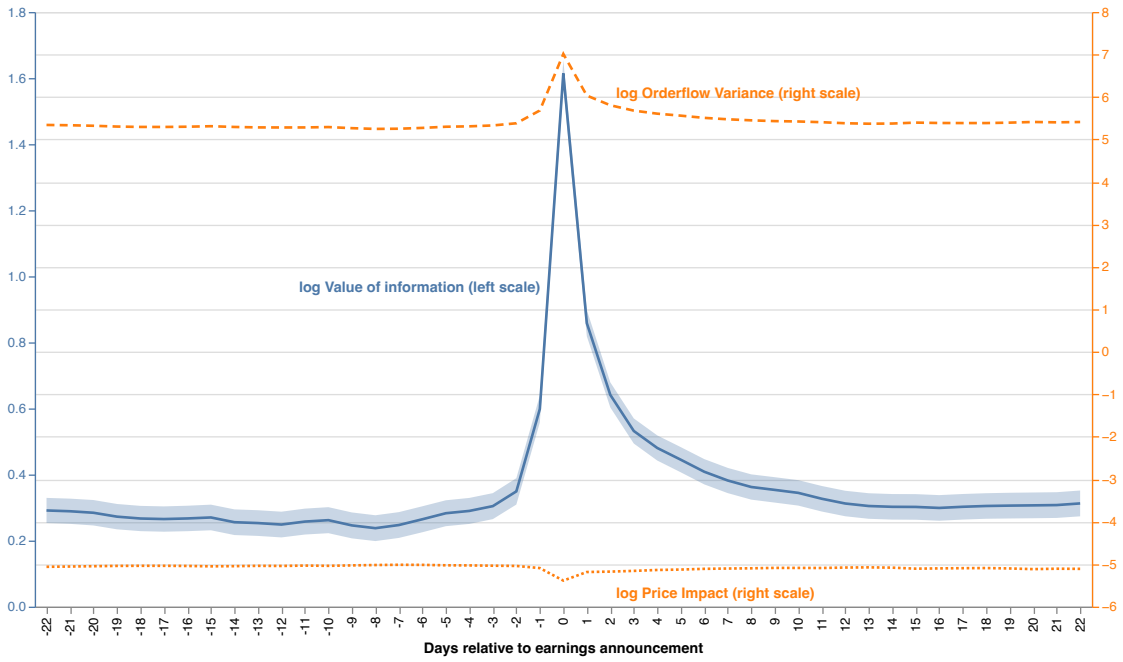
We collect earnings announcement dates from the Compustat quarterly database. Given that firms often report earnings after hours, we follow Engelberg, Mclean, and Pontiff (2018) and define the earnings announcement date as the date in which the trading volume in the firm scaled by market volume is the highest among the official earnings date and the days just before/after it.

Figure 2 plots the value of information averaged across stocks and dates during a 45 trading day window around the earnings announcement date (day 0). The figure shows a dramatic increase in the value of information on days in which earnings are announced. This increase stems primarily from a spike in volatility on these days whereas liquidity appears to be actually mildly improving.

Table 3a provides more formal evidence supporting this conclusion. We regress the log of the value of information for each firm and day during our sample on firm size, book-to-market, momentum, date and stock fixed effects as well as an indicator variable for earning-announcement dates. The coefficient of the earnings indicator is positive and statistically significant in all specifications, indicating that the value of information rises on earning-announcement days. From an economic perspective, this effect is quite large, and implies that the value of information increases more than 3-fold on the mean earnings release ($e^{1.27} \approx 3.57$).

Table 3b sheds further light on the increase in the value of information on earning announcement days. Here, we regress log order flow variance and log price impact on the

Figure 2: Value of information rises when firms report earnings



Notes: Solid line is the daily information value averaged over stocks surrounded by its 95 percent confidence interval, based on time-clustered variance estimates. The value of information is the annualized sum of the product of share price changes times signed share order flow and reported in millions of dollars. The log information value equals log price impact plus log annualized order flow variance, so we also report the mean of log price impact (dotted line) and the mean of log annualized return variance (dashed line).

Table 3: Information about earnings is highly valuable

	log Information value			
	(1)	(2)	(3)	(4)
Earnings	1.38***	1.34***	1.31***	1.27***
	(0.01)	(0.01)	(0.01)	(0.01)
Size	0.70***	0.70***	0.59***	0.62***
	(0.01)	(0.01)	(0.01)	(0.01)
Book-to-market	-0.29***	-0.27***	-0.07***	-0.01
	(0.02)	(0.02)	(0.01)	(0.01)
Momentum	0.29***	0.38***	0.14***	0.21***
	(0.01)	(0.01)	(0.01)	(0.01)
Date (day) FE		Yes		Yes
Stock FE			Yes	Yes
N	13,855,033	13,855,033	13,854,969	13,855,033
R^2	0.53	0.56	0.66	0.69
Within-R^2		0.54	0.12	0.12

(a) Value of information is higher on earnings announcement days

	log Orderflow Variance				log Price impact			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Earnings	1.78***	1.76***	1.67***	1.64***	-0.38***	-0.40***	-0.33***	-0.35***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Size	1.79***	1.81***	1.53***	1.63***	-1.11***	-1.12***	-0.95***	-1.02***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.00)	(0.00)	(0.01)	(0.01)
Book-to-market	-0.34***	-0.32***	-0.17***	-0.04**	0.06***	0.06***	0.10***	0.03***
	(0.03)	(0.03)	(0.02)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)
Momentum	0.45***	0.46***	0.29***	0.27***	-0.16***	-0.08***	-0.15***	-0.06***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.00)
Date (day) FE		Yes		Yes		Yes		Yes
Stock FE			Yes	Yes			Yes	Yes
N	13,855,033	13,855,033	13,854,969	13,855,033	13,855,033	13,855,033	13,854,969	13,855,033
R^2	0.76	0.77	0.84	0.84	0.69	0.72	0.74	0.76
Within-R^2		0.77	0.27	0.25		0.71	0.17	0.16

(b) Order flow volatility is higher and price impact is lower on earnings announcement days

Notes: Panel (a) shows regressions of the log value of information and its components on size, book-to-market, and an earnings indicator. The log information value equals log price impact plus log annualized order flow variance, so in Panel (B) we regress log price impact and log annualized order flow variance on the same variables. Earnings indicates an earnings announcement day. Observations with negative price impact are omitted. Standard errors clustered by date and stock are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

earnings-day indicator and control variables. In all specifications, we find a significant increase in order flow volatility and a decrease in price impact on earning announcement days. These two effects compete with each other, but because the increase in order flow volatility is larger, the value of information increases.

That the value of information rises on earnings announcement days is interesting in its own right. But it can also serve to explain the previously documented fact that investor attention spikes around these events. Ben-Rephael, Da, and Israelsen (2017) and Liu, Peng, and Tang (2019) find that both institutional and retail investor attention rises on earnings days, and that this effect is stronger for institutional investors. This fact is consistent with institutional investors superior ability to gather and exploit private information.

That price impact falls on earnings days is a novel and somewhat suprising finding. A reasonable prior would be that fundamental volatility is greater on earnings days and that market makers would repspond by increasing price impact to protect themselves against adverse selection. But a counterforce is the intensity of noise trading on these days which appears to rise simultaneously. The decline in price impact is also consistent with uninformed noise traders overestimating the novelty of earnings news (Treynor 1971). We are not aware of any other work that has documented this effect and believe it warrants further investigation elsewhere.

Overall, the previous results validate that our proposed measure of the value of information is a useful and informative statistic, and that it yields results that are consistent with prior work.

4 The value of information puzzle

4.1 Economic magnitude of the value of information

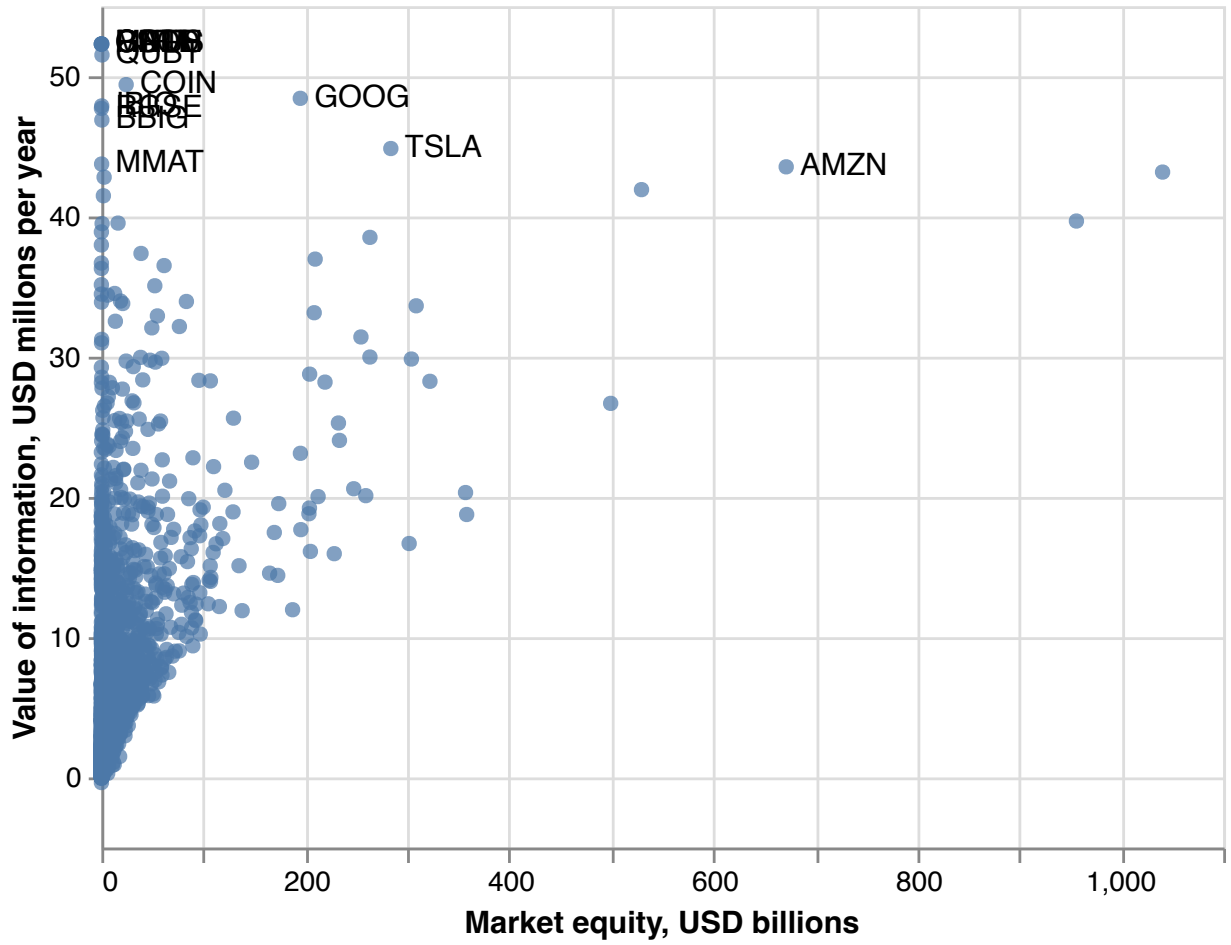
We next step back to ask if the value of information that we estimate is consistent with the demand for information. Recall that from Table 1, the average value of information per stock is about $3.5 million per year. This seems quite small, given that the average market capitalization of a stock in our sample is about $10 billion. But, perhaps the value is considerably larger for some stocks.

Figure 3 shows that indeed for some stocks, the value of information is considerably larger. For example, the value for information on Tesla (TSLA) is about an order of magnitude larger than the average. But even for the most valuable stocks, the value of information is only a small fraction of market capitalization. To see this, in Figure 4 we plot the aggregate value of information as a percentage of aggregate market capitalization.

We estimate that the aggregate value of information is 0.04% of aggregate market capitalization. This aggregate statistic is equivalent to the market cap weighted ratio of the value of information to market capitalization. From Figure 4, we see that the value as percentage of market capitalization is remarkably stable over our sample period, though it tends to spike during turbulent times. Such spikes are partly mechanical, as a drop in stock market prices reduces the denominator in the ratio of the value of information to market capitalization.

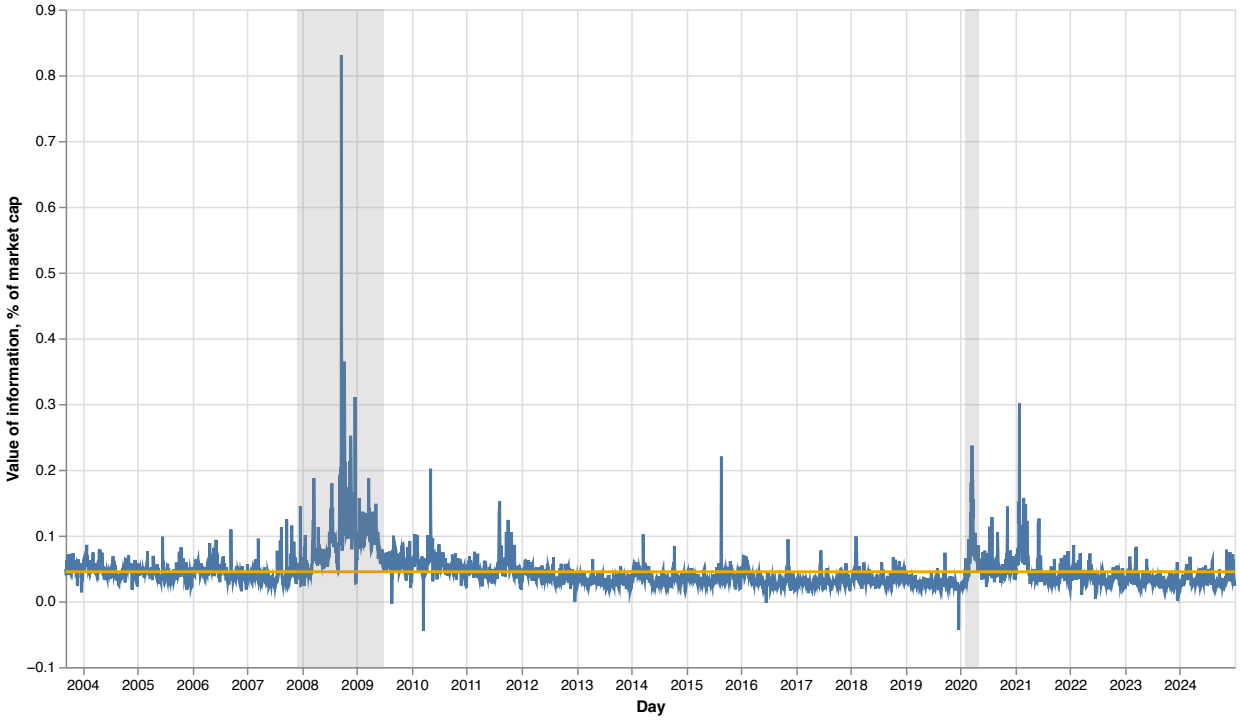
Our estimates imply that the most active investors and asset managers can hope to gain by seeking superior information to forecast stock prices is only 0.04% on average of market capitalization. This estimate is remarkably small compared to the fees that active asset managers charge. For example, French (2008) estimates that investors pay active asset managers about 0.67% of market capitalization in fees each year searching for superior returns. We find it puzzling therefore that investors would pay so much more in fees than

Figure 3: Value of information by stock



Notes: Mean value of information plotted against mean market equity for each specific stock (permno). The value of information is the annualized sum of the product of share price changes times signed share order flow and reported in millions of dollars. We label the 15 most information valuable stocks with their latest ticker symbols.

Figure 4: Value of information as percentage of market capitalization



Notes: Plotted is the daily information value, aggregated over stocks, and divided by the aggregate market capitalization in the prior month. The solid horizontal line is the mean value of information as a percentage of market capitalization (0.04%). Shades indicate recessions.

active asset management can possibly generate.

To further see why this is puzzling, consider the model by Gârleanu and Pedersen (2018), where informed traders, uninformed noise traders, and market markers are joined by asset managers. They show that in a general equilibrium for both assets and asset management, the value of information equals management fees plus the cost of searching for good managers:

$$\text{Value of information} = \text{Management fees} + \text{Search costs} \geq \text{Management fees}. \tag{12}$$

Assuming search costs are nonnegative, this implies that the value of information is at least as large as the fees charged by asset managers. But our estimate of the value of information is much smaller than the fees charged by asset managers.

Note that our puzzle is different from the already puzzling finding in French (2008) that the fees paid to active asset managers do not justify their average outperformance of passive investment strategies. Our findings imply that the fees are high compared the most that active asset managers could *possibly* generate.

In Appendix A.1 we show that this puzzling fact is robust to a battery of robustness tests. First, we study the sensitivity of our estimates to how we sign trades as buys or sells. In Figure A.1 we report the mean value of information across days and stocks for three different algorithms for signing trades as buys or sells: Chakrabarty et al. (2007, CLNV), Lee and Ready (1991, LR), and Ellis, Michaely, and O'Hara (2000, EMO). None of these have a major effect on the mean value of information. Second, in our baseline analysis we omit penny stocks with less than \$5 in price. The figure further shows that including such stocks would reduce the mean value of information. This is intuitive as penny stocks have little order flow and therefore little information value. In Figure A.2 we show that our estimates are also robust to variation in the frequency of observation.

4.2 Potential resolutions

What can possibly explain the fact that the value of information that we estimate is so small compared to the fees that active asset managers charge? We next consider several potential resolutions.

4.2.1 Decline in fees

One potential concern is that French (2008)'s estimates of aggregate fees as a percentage of market capitalization are outdated. If active asset management fees have declined considerably since 2006, when French's sample ends, this could resolve the puzzle because our sample period is September 2003 to December 2024.

Table 4: Cost of Active vs. Passive Equity Investing (% of AUM)

	Morningstar (MF + ETF)			**ICI Mutual Funds**			**ICI ETFs**		
Year	Active	Passive	Diff.	Active	Passive	Diff.	Active	Passive	Diff.
2000	1.03	0.30	0.73	1.06	0.27	0.79			
2010	0.82	0.25	0.57	0.96	0.16	0.80			
2016	0.71	0.17	0.54				0.88	0.22	0.66
2018	0.67	0.14	0.53				0.79	0.19	0.60
2020	0.64	0.12	0.52	0.71	0.06	0.65	0.70	0.18	0.52
2022	0.61	0.11	0.50				0.44	0.16	0.28
2024	0.59	0.11	0.48	0.64	0.05	0.59	0.44	0.14	0.30
Mean	*0.72*	*0.17*	*0.55*	*0.84*	*0.14*	*0.71*	*0.65*	*0.18*	*0.47*

Notes: All figures are asset-weighted average expense ratios for equity funds expressed as a percentage of assets under management. "Active" refers to actively managed funds; "Passive" refers to index funds. "Diff." is the difference (Active − Passive). Morningstar data from their 2024 U.S. Fund Fee Study combine mutual funds and ETFs; they report that ETF fees are less than half of mutual fund fees on average. ICI data from the 2025 Investment Company Fact Book, Tables 6.4 (mutual funds) and 6.5 (ETFs), report mutual funds and ETFs separately. For comparison, French (2008) estimates aggregate fees of 0.67% of market capitalization for the period 1980–2006, which includes all costs of active investing relative to a passive benchmark.

In Table 4, we compare the expense ratios of active versus passive (index) fund expense

ratios over time, focusing on equity funds. Similar to French (2008), we use the difference between active and passive fees to measure the additional cost investors pay for active management. This difference has narrowed over time but remains substantial: in 2024, actively managed equity mutual funds charged 0.64% versus 0.05% for index funds (ICI data), a difference of 59 basis points. For ETFs, actively managed equity funds charged 0.44% versus 0.14% for index ETFs, a difference of 30 basis points. Evidently, the fee compression observed since French's sample period does not resolve the puzzle.

4.2.2 Additional benefits of asset management

One might argue that the fees investors pay to asset managers compensate not just for private information, but for a bundle of services: professional portfolio construction, diversification, factor or style exposures, liquidity transformation, tax management, custodial services, and convenience. If investors value these non-informational services, perhaps the gap between our estimate of the value of information and aggregate fees is not puzzling after all.

However, this argument does not resolve the puzzle. Passive index funds provide essentially the same bundle of services—diversification, factor exposure, liquidity, custody, and convenience—at dramatically lower cost. As documented in Table 4, the fee differential between active and passive equity funds is substantial: 59 basis points for mutual funds and 30 basis points for ETFs in 2024. This differential represents the incremental cost investors pay specifically for active management—that is, for the promise of superior returns through private information or stock-picking skill.

Our estimates suggest that the aggregate value of information available in the market is only about 0.04% of market capitalization. This is far smaller than the active-passive fee differential, let alone the total fees charged by active managers. Thus, even if we grant that some portion of active management fees compensates for non-informational services, the

portion attributable to active management remains puzzlingly large relative to the value such activity can generate.

4.2.3 Perhaps being active on some days is worth it?

One might argue that the puzzle could be resolved if active managers strategically concentrate their information acquisition on days when the value of information is particularly high. Earnings announcement days are a natural candidate, as we documented in Section 3.5 that the value of information increases more than 3-fold on these days.

However, even this concentration strategy cannot bridge the gap. On earnings days, the value of information rises from 0.04% to roughly $0.04 \times 3.57 \approx 0.14\%$ of market capitalization. This is still considerably less than the 0.67% in fees that investors pay searching for superior returns.

Moreover, earnings days represent only a tiny fraction of trading days—about 1% of stock-days in our sample. A selective strategy that captures information value only on these high-value days would therefore earn approximately $0.01 \times 3.57 \times 0.04 \approx 0.0014\%$ of market capitalization per year—far less than the 0.04% one would earn by being active on all days. This calculation underscores that even on the most information-rich days in a firm's calendar, being active does not come close to justifying the fees charged by active asset managers.

4.2.4 Noncompetitive market makers

Our first main assumption in deriving Proposition 1 is that the total expected gain to informed investors equals the total expected losses to noise traders (Assumption 1). This equality holds if market makers are competitive and earn no rents in equilibrium. It is therefore natural to consider relaxing this assumption. Doing so, however, would deepen the puzzle.

Suppose that instead, market markers have some market power and charge a markup. In such a market, market makers receive the same amount of losses from noise traders, but pass only part of them to informed investors. As a result, aggregate losses by noise traders provide an upper bound on the value of information.

4.2.5 Noncompetitive asset managers

The aggregate fees estimated by French simply measure aggregate fees paid by investors, and do not assume anything about the competitiveness of asset management. Similarly, the fund specific expense ratios we use in Table 4 measure the fees paid by investors to the average fund, and do not assume anything about their competitiveness. The basic intuition, formalized in Equation (12), suggests that the value of information is bounded from below by the fees paid by investors to the asset managers in equilibrium, regardless of the competitiveness of the asset management. This intuition is clearly violated by our estimates.

4.2.6 Risk averse informed investors

In Kyle-Back models, informed investors are risk neutral. By contrast, informed investors are risk averse in Subrahmanyam (1991)'s extension of the Kyle model, and in noisy rational expectations models in the tradition of Grossman and Stiglitz (1980). How much would risk averse informed investors be willing to pay in this environment? Well, the aggregate losses of the noise traders would still measure the expected gain in monetary terms from superior information. But due to risk aversion and Jensen's inequality, they would be willing to pay less than the expected gain. Thus, higher risk aversion would tend to lower the value of information and further deepen the puzzle.

4.2.7 Limited arbitrage and risk averse market makers

In the Kyle-Back framework, market makers are unconstrained in their ability to arbitrage the asset. In reality, market makers are constrained by the need to balance their inventory risk and the need to earn a profit.

Kondor and Vayanos (2019) study the dynamic interactions between liquidity providers' capital, the liquidity that these agents provide to other participants, and assets' risk premia. They show that in a framework with minimal frictions, in particular, no asymmetric information or leverage constraints, risk averse market markers (arbitrageurs) would set positive price impact coefficients (λ) that decrease with their wealth. That is, even if no asymmetric information is present, price changes could still covary with order flow, and we could misattribute this to the value of information.

We are aware of no theoretical work that combines their model with asymmetric information, but it is plausible that in such an extension of their model, price impact coefficient would compensate market makers for both adverse selection and risk aversion. In this case too, we would underestimate the value of information, and the puzzle would deepen.

4.2.8 Partial information and multiple informed investors

In Kyle (1985) and Back (1992) framework, a monopolistic informed investor learns the fundamental value of the asset perfectly, without noise. It is arguably more realistic that the informed receive only a noisy signal about the asset's fundamental value. Moreover, multiple investors can simultaneously acquire information about the asset. Would such features change the interpretation of our estimates?

Back, Cao, and Willard (2000) consider a more general model in which multiple investors receive partial and potentially correlated signals about an asset's fundamental value. Although this model is harder to solve and generates different dynamics of trade and prices, it generates the same expression for the aggregate losses of the uninformed

noise traders and for the aggregate value of information to all informed investors combined (Back, Cao, and Willard 2000, Corollary 3).

Intuitively, even with multiple informed investors and partial signals, the gains to the informed come entirely from the losses of the noise traders. Therefore, our approach of focusing on the losses of the noise traders allows us to abstract from these specific model features.

4.2.9 Discrete time and relaxing orthogonal informed trading Assumption 2

Assumption 2 holds exactly in continuous-time Kyle-Back models because informed trading is locally deterministic. But what if we consider a more general discrete-period setting where this assumption may be violated? Here we show that even in such cases, our measure provides an *upper bound* on the value of information.

Consider a discrete-time analog where trading occurs over N periods. Let $\Delta P_n = P_n - P_{n-1}$, $\Delta Y_n = Y_n - Y_{n-1}$, $\Delta X_n = X_n - X_{n-1}$, and $\Delta Z_n = Z_n - Z_{n-1}$ denote the changes in prices, order flow, informed trading, and noise trading in period n. Assume that price changes are linear in order flow:

$$\Delta P_n = \lambda_n \Delta Y_n, \tag{13}$$

for some positive price impact coefficient $\lambda_n > 0$. We further assume that noise trading innovations are uncorrelated with informed trading decisions:

$$E[\Delta Z_n \Delta X_n] = 0. \tag{14}$$

This assumption is natural: noise traders' liquidity needs are exogenous and independent of informed traders' strategies, which depend on private information about fundamentals.

The expected profit to informed traders over all periods can be written as:

$$\begin{aligned}\Omega &= E\sum_{n=1}^{N}\Delta P_n \Delta Z_n \\ &= E\sum_{n=1}^{N}\Delta P_n(\Delta Y_n - \Delta X_n) \\ &= E\sum_{n=1}^{N}\Delta P_n \Delta Y_n - E\sum_{n=1}^{N}\Delta P_n \Delta X_n. \end{aligned} \tag{15}$$

The first term is precisely what we measure empirically. The second term captures the extent to which Assumption 2 is violated in discrete time.

Using the linear pricing rule (13), the second term becomes:

$$\begin{aligned}E\sum_{n=1}^{N}\Delta P_n \Delta X_n &= E\sum_{n=1}^{N}\lambda_n \Delta Y_n \Delta X_n \\ &= E\sum_{n=1}^{N}\lambda_n(\Delta X_n + \Delta Z_n)\Delta X_n \\ &= E\sum_{n=1}^{N}\lambda_n(\Delta X_n)^2, \end{aligned} \tag{16}$$

where the last equality follows from the orthogonality condition (14). Since $\lambda_n > 0$, this term is non-negative. Therefore:

$$\Omega = E\sum_{n=1}^{N}\Delta P_n \Delta Y_n - E\sum_{n=1}^{N}\lambda_n(\Delta X_n)^2 \leq E\sum_{n=1}^{N}\Delta P_n \Delta Y_n. \tag{17}$$

This establishes that our empirical measure $E\sum \Delta P_n \Delta Y_n$ provides an upper bound on the true value of information even when Assumption 2 does not hold exactly. If anything, violating this assumption would make the puzzle even deeper: the true value of information would be even smaller than our already modest estimates suggest.

Moreover, the "error term" in (16) vanishes as the time step shrinks. In Kyle-style models, informed trading increments ΔX_n are of order Δt (smooth), so $(\Delta X_n)^2$ is of order $(\Delta t)^2$, which becomes negligible relative to $\Delta P_n \Delta Y_n$ (of order Δt) as $\Delta t \to 0$. This provides additional intuition for why Assumption 2 holds exactly in continuous time and approximately in high-frequency discrete time.

4.2.10 Jumps in prices

The continuous-time Kyle-Back framework assumes that prices follow diffusion processes without jumps. In reality, asset prices can exhibit discontinuous jumps, particularly around news events or market stress. Importantly, the discrete-time analysis developed in Section 4.2.9 implicitly accommodates such jumps.

In our discrete-period extension in Section 4.2.9, we only require that price changes satisfy $\Delta P_n = \lambda_n \Delta Y_n$ for some positive λ_n. This relation places no restriction on the magnitude or continuity of price changes between periods. Prices may jump discontinuously from P_{n-1} to P_n, and the upper bound in (17) still applies. Therefore, the result—that $E \sum \Delta P_n \Delta Y_n$ provides an upper bound on the value of information—extends naturally to settings with price jumps.

This generalization is empirically relevant because our high-frequency estimates aggregate price changes and order flow over discrete time intervals (e.g., one-minute windows). Within each interval, prices may exhibit jumps due to news arrivals or large trades. The discrete-time framework ensures that our estimates remain valid upper bounds on the value of information even in the presence of such discontinuities.

4.2.11 Behavioral explanations

Our discussion above suggests that solving the puzzle requires an additional source of revenue for market makers that would compensate them for losses to informed traders,

even if price impact is as low as our estimates suggest. Treynor (1971) writes about a third type of trader that is potentially missing from the Kyle-Back framework—"transactors acting on information which they believe has not yet been fully discounted in the market price but which in fact has." Treynor (1971) explains,

> This is where the third kind of transactor comes in: From the market maker's point of view, his effect is identical to the liquidity-motivated transactor's. The market maker naturally welcomes the cooperation of wire houses and information services like the Wall Street Journal that broadcast information already fully discounted since many investors are easily persuaded to transact based on that information, hence enable the market maker to maintain substantially smaller spreads than would be possible without their trading activity.

While Treynor (1971) thought of such traders as identical in their effect on the market to noise traders, one may imagine their misunderstanding of their position in the information supply chain, may lead them to trade in the wrong direction on average. In this case, such behavioral investors would lose more than the bid-ask spread in each trade.

Large systematic trading mistakes are a potential explanation for our puzzling estimates. In such an environment, the aggregate losses that are shared by market makers and informed traders are composed of both trading costs which we measure and trading mistakes which we do not.

Barber et al. (2009) quantify the losses of individual investors using complete transaction data from Taiwan. They find that individual investors lose approximately 2.3% of aggregate market capitalization annually between 1995 and 1999. Of this, gross trading losses—the pure stock-picking losses from buying stocks that underperform stocks sold—amount to 0.6% of market capitalization annually, with the remainder attributable to commissions, transaction taxes, and market-timing losses. More recent evidence from US markets documents similar patterns. Barber et al. (2022) find that intense buying by Robin-

hood users forecasts negative returns, with average 20-day abnormal returns of -4.7% for the top stocks purchased each day. Barber, Lin, and Odean (2024) show that retail purchases are concentrated in attention-grabbing stocks that subsequently underperform, with long-short strategies based on retail order imbalance yielding -14.8% annualized returns among stocks with heavy retail trading. While these studies document systematic trading mistakes by US retail investors, they do not provide aggregate estimates scaled by market capitalization. Whether such losses are large enough to explain our puzzling estimates remains an open question.

5 Conclusion

We develop a simple, general closed-form expression for the value of information in financial markets. Under mild assumptions satisfied by a broad class of Kyle-Back models—competitive market making and orthogonal informed trading—the value of information equals the expected quadratic covariation between price changes and order flow. This expression is easy to estimate using high-frequency data and has straightforward economic units.

Using intraday TAQ data on US equities over September 2003 to December 2024, we estimate that the average annualized value of information is approximately $3.5 million per stock. In aggregate, the value of information amounts to only 0.04% of market capitalization. Our estimates pass several validation tests: the value of information rises during turbulent periods, is higher for large, growth, and momentum stocks, and increases more than 3-fold on earnings announcement days.

These findings present a puzzle. The aggregate value of information we estimate is considerably smaller than the approximately 0.67% of market capitalization that investors pay in fees to active asset managers each year (French 2008). Even accounting for fee compression over time and focusing on the active-passive fee differential, this gap remains

substantial. We consider several potential resolutions—including non-competitive market makers, risk aversion, partial information, and behavioral trading mistakes—but find that most would either deepen the puzzle or leave it unresolved.

One interpretation of our findings is that the equity market is highly informationally efficient, leaving little value for informed traders to extract. An alternative interpretation is that a significant portion of active management fees compensates for services unrelated to information, or that behavioral investors systematically lose more than the trading costs we measure. Understanding which forces drive this wedge between the value of information and the cost of seeking it remains an important question for future research.

References

Ai, Hengjie, Ravi Bansal, Hongye Guo, and Amir Yaron. 2026. "Identifying Preference for Early Resolution from Asset Prices." *American Economic Review.*

Amihud, Yakov. 2002. "Illiquidity and Stock Returns: Cross-Section and Time-Series Effects." *Journal of Financial Markets* 5, no. 1 (January): 31–56. ISSN: 1386-4181.

Back, Kerry. 1992. "Insider Trading in Continuous Time." *Review of Financial Studies* 5 (3): 387–409. ISSN: 0893-9454.

Back, Kerry, C. Henry Cao, and Gregory A. Willard. 2000. "Imperfect Competition among Informed Traders." *The Journal of Finance* 55, no. 5 (October): 2117–2155. ISSN: 00221082.

Back, Kerry, Kevin Crotty, and Tao Li. 2018. "Identifying Information Asymmetry in Securities Markets." *The Review of Financial Studies* 31, no. 6 (June): 2277–2325. ISSN: 0893-9454.

Bai, Jennie, Thomas Philippon, and Alexi Savov. 2016. "Have Financial Markets Become More Informative?" *Journal of Financial Economics* 122, no. 3 (December): 625–654. ISSN: 0304-405X.

Barber, Brad M., Xing Huang, Terrance Odean, and Christopher Schwarz. 2022. "Attention-Induced Trading and Returns: Evidence from Robinhood Users." *The Journal of Finance* 77 (6): 3141–3190. ISSN: 1540-6261.

Barber, Brad M., Yi-Tsung Lee, Yu-Jane Liu, and Terrance Odean. 2009. "Just How Much Do Individual Investors Lose by Trading?" *The Review of Financial Studies* 22, no. 2 (February): 609–632. ISSN: 0893-9454.

Barber, Brad M., Shengle Lin, and Terrance Odean. 2024. "Resolving a Paradox: Retail Trades Positively Predict Returns but Are Not Profitable." *Journal of Financial and Quantitative Analysis* 59, no. 6 (September): 2547–2581. ISSN: 0022-1090, 1756-6916.

Ben-Rephael, Azi, Zhi Da, and Ryan D. Israelsen. 2017. "It Depends on Where You Search: Institutional Investor Attention and Underreaction to News." *The Review of Financial Studies* 30, no. 9 (September): 3009–3047. ISSN: 0893-9454.

Bogousslavsky, Vincent, Vyacheslav Fos, and Dmitriy Muravyev. 2024. "Informed Trading Intensity." *The Journal of Finance* 79 (2): 903–948. ISSN: 1540-6261.

Bond, Philip, Alex Edmans, and Itay Goldstein. 2012. "The Real Effects of Financial Markets." *Annual Review of Financial Economics* 4, no. 1 (October): 339–360. ISSN: 1941-1367.

Breon-Drish, Bradyn. 2015. "On Existence and Uniqueness of Equilibrium in a Class of Noisy Rational Expectations Models." *Review of Economic Studies*, rdv012.

Brogaard, Jonathan, Matthew C. Ringgenberg, and David Sovich. 2019. "The Economic Impact of Index Investing." *Review of Financial Studies* 32 (9): 3461–3499. ISSN: 0893-9454.

Chakrabarty, Bidisha, Bingguang Li, Vanthuan Nguyen, and Robert A. Van Ness. 2007. "Trade Classification Algorithms for Electronic Communications Network Trades." *Journal of Banking & Finance* 31 (12): 3806–3821. ISSN: 0378-4266.

Collin-Dufresne, Pierre, and Vyacheslav Fos. 2016. "Insider Trading, Stochastic Liquidity, and Equilibrium Prices." *Econometrica* 84 (4): 1441–1475. ISSN: 1468-0262.

Croce, Mariano M, Tatyana Marchuk, and Christian Schlag. 2023. "The Leading Premium." *The Review of Financial Studies* 36, no. 8 (August): 2997–3033. ISSN: 0893-9454.

Dávila, Eduardo, and Cecilia Parlatore. 2021. "Trading Costs and Informational Efficiency." *Journal of Finance* 76 (3): 1471–1539. ISSN: 1540-6261.

———. 2023. "Volatility and Informativeness." *Journal of Financial Economics* 147 (3): 550–572.

———. 2025. "Identifying Price Informativeness." *The Review of Financial Studies* (July): hhaf051. ISSN: 0893-9454.

Ellis, Katrina, Roni Michaely, and Maureen O'Hara. 2000. "The Accuracy of Trade Classification Rules: Evidence from Nasdaq." *Journal of Financial and Quantitative Analysis* 35 (4): 529–551. ISSN: 0022-1090.

Engelberg, Joseph, R. David Mclean, and Jeffrey Pontiff. 2018. "Anomalies and News." *Journal of Finance* 73 (5): 1971–2001. ISSN: 1540-6261.

Epstein, Larry G, and Stanley E Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica*, 937–969.

Epstein, Larry G., Emmanuel Farhi, and Tomasz Strzalecki. 2014. "How Much Would You Pay to Resolve Long-Run Risk?" *American Economic Review* 104 (9): 2680–2697.

Fama, Eugene F. 1970. "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance* 25, no. 2 (May): 383–417.

Farboodi, Maryam, Adrien Matray, Laura Veldkamp, and Venky Venkateswaran. 2021. "Where Has All the Data Gone?" *The Review of Financial Studies,* no. hhab110 (October). ISSN: 0893-9454.

Farboodi, Maryam, Dhruv Singal, Laura Veldkamp, and Venky Venkateswaran. 2025. "Valuing Financial Data." *The Review of Financial Studies* 38, no. 3 (March): 938–980. ISSN: 0893-9454.

French, Kenneth R. 2008. "Presidential Address: The Cost of Active Investing." *Journal of Finance* LXIII (4): 1537–1573.

Gârleanu, Nicolae, and Lasse Heje Pedersen. 2018. "Efficiently Inefficient Markets for Assets and Asset Management." *Journal of Finance* 73 (4): 1663–1712. ISSN: 1540-6261.

Goldstein, Itay. 2023. "Information in Financial Markets and Its Real Effects*." *Review of Finance* 27, no. 1 (February): 1–32. ISSN: 1572-3097.

Grossman, Sanford J., and Joseph E. Stiglitz. 1980. "On the Impossibility of Informationally Efficient Markets." *American Economic Review* 70 (3): 393–408. ISSN: 00028282.

Hasbrouck, Joel. 2009. "Trading Costs and Returns for U.S. Equities: Estimating Effective Costs from Daily Data." *Journal of Finance* 64 (3): 1445–1477. ISSN: 1540-6261.

Holden, Craig W., and Stacey Jacobsen. 2014. "Liquidity Measurement Problems in Fast, Competitive Markets: Expensive and Cheap Solutions." *Journal of Finance* 69 (4): 1747–1785. ISSN: 1540-6261.

Kacperczyk, Marcin, Savitar Sundaresan, and Tianyu Wang. 2021. "Do Foreign Institutional Investors Improve Price Efficiency?" *Review of Financial Studies* 34, no. 3 (March): 1317–1367. ISSN: 0893-9454.

Kadan, Ohad, and Asaf Manela. 2019. "Estimating the Value of Information." *Review of Financial Studies* 32 (3): 951–991. ISSN: 0893-9454.

———. 2025. "Liquidity and the Strategic Value of Information." *Review of Finance* 29, no. 1 (January): 1–32. ISSN: 1572-3097, 1573-692X.

Kondor, Péter, and Dimitri Vayanos. 2019. "Liquidity Risk and the Dynamics of Arbitrage Capital." *The Journal of Finance* 74 (3): 1139–1173. ISSN: 1540-6261.

Kyle, Albert S. 1985. "Continuous Auctions and Insider Trading." *Econometrica* 53 (6): 1315–1335. ISSN: 00129682.

Kyle, Albert S., and Anna A. Obizhaeva. 2016. "Market Microstructure Invariance: Empirical Hypotheses." *Econometrica* 84 (4): 1345–1404.

———. 2018. "The Market Impact Puzzle." *Anna A., The Market Impact Puzzle* (*February 4, 2018*).

———. 2019. "Market Microstructure Invariance: A Dynamic Equilibrium Model." *Anna A., Market Microstructure Invariance: A Dynamic Equilibrium Model* (*January 31, 2019*).

Kyle, Albert S., Anna A. Obizhaeva, and Yajun Wang. 2018. "Smooth Trading with Overconfidence and Market Power." *Review of Economic Studies* 85 (1): 611–662.

Lee, Charles M. C., and Mark J. Ready. 1991. "Inferring Trade Direction from Intraday Data." *Journal of Finance* 46 (2): 733–746. ISSN: 1540-6261.

Liu, Hongqi, Lin Peng, and Yi Tang. 2019. *Retail Attention, Institutional Attention.* SSRN Scholarly Paper ID 3370683. Rochester, NY: Social Science Research Network, April.

Malamud, Semyon. 2015. *Noisy Arrow-Debreu Equilibria.* Working Paper.

Malamud, Semyon, and Marzena Rostek. 2017. "Decentralized Exchange." *American Economic Review* 107, no. 11 (November): 3320–3362. ISSN: 0002-8282.

Obizhaeva, Anna A. 2011. "Selection Bias in Liquidity Estimates." *University of Maryland Working.*

Savov, Alexi. 2014. "The Price of Skill: Performance Evaluation by Households." *Journal of Financial Economics* 112 (2): 213–231. ISSN: 0304-405X.

Subrahmanyam, Avanidhar. 1991. "Risk Aversion, Market Liquidity, and Price Efficiency." *The Review of Financial Studies* 4, no. 3 (July): 417–441. ISSN: 0893-9454.

Treynor, Jack. 1971. "The Only Game in Town." *Financial Analysts Journal* 27 (2): 12–22. ISSN: 0015-198X.

Veldkamp, Laura. 2011. *Information Choice in Macroeconomics and Finance.* Princeton, NJ: Princeton University Press.

A Appendix

A.1 Robustness

Here we include some additional robustness tests.

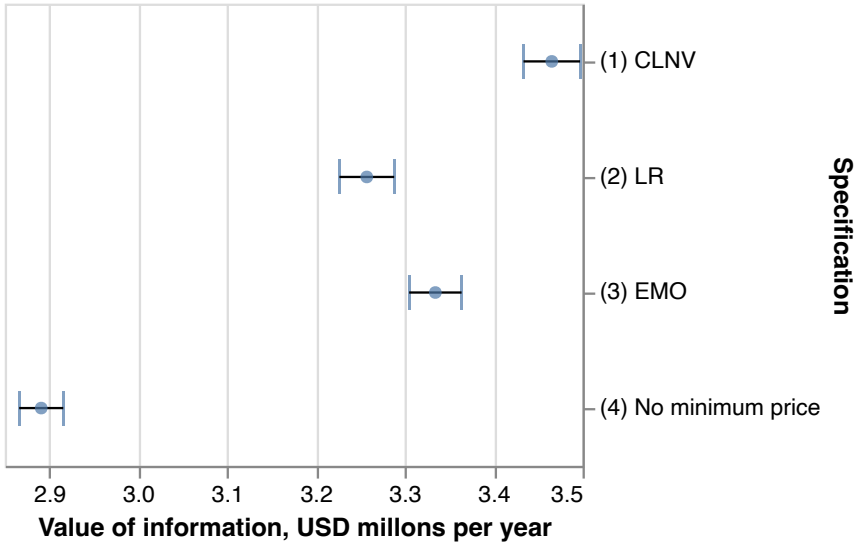
To study the sensitivity of our estimates to how we estimate price impact, in Figure A.1 we report the mean value of information across days and stocks for three different algorithms for signing trades as buys or sells: Chakrabarty et al. (2007, CLNV), Lee and Ready (1991, LR), and Ellis, Michaely, and O'Hara (2000, EMO). None of these have a major effect on the mean value of information.

Our baseline analysis omits penny stocks with less than $5 in price. The figure further shows that including such stocks would reduce the mean value of information. This is intuitive as penny stocks likely have large price impact coefficients.

In Figure A.2 we show that our estimates of the value of information rise only slightly with the frequency of observation. As we show in Section 4.2.9, the error term in (16) vanishes as the time step shrinks. Therefore, higher frequency estimates provide a better approximation to the value of information, while lower frequency estimates are biased upwards. The figure shows that even at 30-minute frequency we estimate a modest upper bound on the value of information.

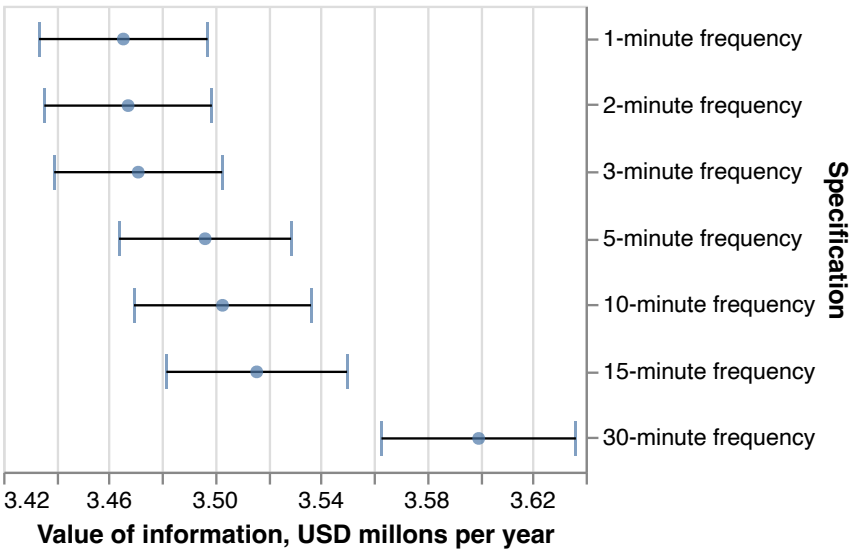
Figure A.3 reports similar robustness tests, but this time for the regression analysis reported in the baseline regression tables in the paper (e.g. Table 3). Again, we find that our conclusions are robust to these alternative specifications.

Figure A.1: Robustness: Mean value of information



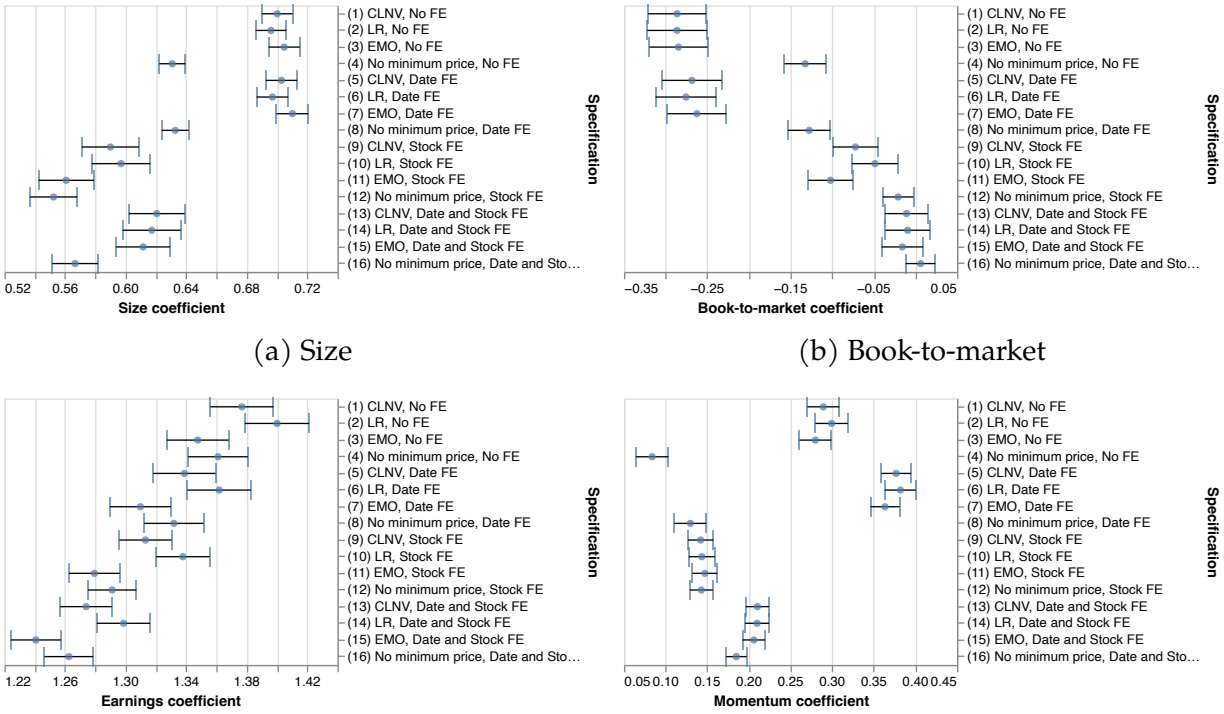
Notes: Reported is the mean value of information across days and stocks for three different algorithms for signing trades as buys or sells: Chakrabarty et al. (2007, CLNV), Lee and Ready (1991, LR), and Ellis, Michaely, and O'Hara (2000, EMO). No minimum price shows the effect of including stocks with nominal price below $5. Idiosyncratic variance shows the effect of subtracting market variance from stock-specific variance to measure the values of idiosyncratic information. The value of information is the annualized sum of the product of share price changes times signed share order flow and reported in millions of dollars. All variables are 1% winsorized.

Figure A.2: Robustness to the frequency of observation



Notes: Reported is the mean value of information across days and stocks for various choices for the frequency of observation. The value of information is the annualized sum of the product of share price changes times signed share order flow and reported in millions of dollars. All variables are 1% winsorized.

Figure A.3: Robustness: Value of information and stock characteristics



(a) Size

(b) Book-to-market

(c) Earnings

(d) Momentum

Notes: Reported are coefficients from regressions of the log value of information on stock characteristics. We evaluate several fixed effect specifications and three different algorithms for signing trades as buys or sells: Chakrabarty et al. (2007, CLNV), Lee and Ready (1991, LR), and Ellis, Michaely, and O'Hara (2000, EMO). No minimum price shows the effect of including stocks with nominal price below \$5. Size is log market equity over the previous month. Momentum is the return over the prior 2–12 months. Book-to-market is book equity as of last June divided by market equity as of last December (Fama-French conventions). Observations with negative price impact are omitted. Standard errors clustered by date and stock are used to construct the 95% errors bars.