Detecting Informed Trade by Corporate Insiders

Abstract

Detecting informed trade by corporate insiders is costly and is the subject of significant regulatory and market scrutiny. We introduce a mixture model that leverages the cross-section of insiders' past returns to infer which insiders are more likely to engage in informed trade. The estimation explicitly accounts for the noisiness of insiders' performance histories. Out-of-sample returns are higher for stocks traded by insiders identified as more likely to use information, and prices reflect this information faster over the last decade. The model for insiders implies a person-specific mixture distribution that can be used to classify whether *any* disclosed trade is informed.

November 18, 2024

1. Introduction

Managers, directors, and other firm insiders are, by definition, endowed with private information about their firms. Concerns about insiders' ability to profit from trade on this information go back to the first public company, when the Dutch East India Company banned directors from trading in the stock (Koppell, 2011). Since that time, the costs, benefits, and appropriate regulations of insiders' trades have been extensively debated, not only by policymakers and market participants, but also in the academic fields of economics, finance, accounting, and law. Insider trading has been argued to benefit, for example, production decisions, information provision, and even efficient compensation design through better price efficiency. On the other hand, insider trades adversely select outside investors, which may hinder market liquidity and even market participation.

Of course, not all trade by corporate insiders is necessarily informed. Many insiders hold a substantial fraction of wealth in their firm's equity, so they may trade out of liquidity or diversification concerns. Moreover, there is likely heterogeneity in insiders' propensity to profit from their informational advantage due to ethical or legal considerations. Regulators tasked with enforcing prohibitions on trading on material and nonpublic information as well as scholars interested in studying the economic consequences of informed insider trade face an enormous challenge: the extent of informed versus uninformed trade by corporate insiders is not directly observable.

For publicly traded firms in the U.S., we do observe the return histories of insiders' trades. These histories provide a very noisy signal about an insider's use of information. The signal-to-noise ratio is low not only because some trade is liquidity-motivated, but also because individual stock returns are quite volatile and the historical record for many corporate insiders is short. Indeed, notable existing proxies for informed trade by corporate insiders do not use past returns at all, instead resorting to the persistence in calendar timing of trade (Cohen, Malloy, and Pomorski, 2012) or the consistency of trading direction (Akbas, Jiang, and Koch, 2020).

In this paper, we embrace the idea that signals and noise differ across insiders. Importantly, we take advantage of the fact that there is additional information in the cross-section of signals. We apply mixture model methods that leverage the cross-section of insiders' return histories to (1) infer which insiders are more likely to engage in informed trade and (2) classify which individual trades are more likely to be informationally motivated. The method is specifically designed to reduce the effect of noise on inference about the extent of information embedded in corporate insider trades.[1] In a sample of disclosed stock trades by corporate insiders from 1985 to 2022, about 30% of all insiders fall into the distribution of those trading on private information. On average, the insiders in this distribution earn 3.6% abnormal returns over the next month compared to those not trading on private information, who earn 0% on average (by construction).

Using the estimated mixture model parameters and a realized average abnormal return and standard error for each individual insider, we estimate a conditional probability that a given insider makes informed trades as well as a conditional expected average abnormal return. The mixture model essentially functions as a noise reduction method, where there are multiple ways for an insider to receive the same estimated conditional probability of informed trading. Intuitively, the econometrician should update more strongly that an insider trades on private information if the insider's average return is higher. But estimation noise due to volatile returns or a short trading history should affect this inference. Consider two insiders who both have average abnormal returns of 1%, but the standard error of the first insider's average is 1% while the second insider's standard error is 5%. It is much more likely that the first insider trades on private information than the second insider. Put differently, it is more likely that the second insider's 1% average return occurred by chance than for the first insider. The conditional probability in the mixture model formally quantifies this intuition.

To validate our model's ability to identify traders who are more likely to trade on private information, we test that our estimates have economic content out-of-sample. At any point in time, we can estimate each insider's probability of trading on information using their full trade history up until that point. We conduct out-of-sample exercises testing whether the buys and sells of insiders with different ex-ante conditional expectations predict differences in stock returns. The difference in future returns for stocks with buying activity by top-quintile insiders relative to bottom-quintile insiders is 82 basis points per month or almost 10% per year. Future returns for stocks with selling activity by top-quintile insiders are 46 basis points lower per month (i.e., 5.5% annually) than returns of stocks sold by bottom-quintile insiders. Our results are, therefore,

[1]Mixture models have been used to assess the extent of repeatable performance of various financial market participants, including hedge funds (Chen, Cliff, and Zhao, 2017), mutual funds (Harvey and Liu, 2018), security analysts (Crane and Crotty, 2020), and social media analysts (Kakhbod, Kazempour, Livdan, and Schürhoff, 2023; Dim, 2023).

broadly consistent with prior literature, which finds that insider purchases are more likely to be informed than insider sales. We find that both buys and sells of the most informed traders predict future returns, though the magnitude of the return predictability is substantially larger for insider purchases.

We show that the market has learned about heterogeneity in insiders' propensities to utilize private information, and prices have become more efficient with respect to this information over time. The time horizon of return predictability has changed dramatically over the last thirty years. In the last decade or so, returns are no longer predictable using monthly rebalancing. The information content of insider trading is still present over this time period, but one must form portfolios much closer to the trade's disclosure than the monthly frequency. Interestingly, this apparent faster convergence to market efficiency follows substantial increases in information acquisition by sophisticated investors. We document a number of hedge funds that dramatically increase the acquisition of insider trade disclosures from the SEC website prior to the decline in profitability at the monthly horizon. This increase coincides with the circulation of the influential Cohen et al. (2012) paper on opportunistic insider trading. Our results thus echo those of McLean and Pontiff (2016) in that markets converge more quickly to efficient prices over time through a combination of academic research and information acquisition by sophisticated institutions.

A number of important papers identify cross-sectional differences in which insiders engage in informed trading by conditioning on ex-ante trading patterns predicted to be correlated with using private information. For instance, Cohen, Malloy, and Pomorski (2012) identifies insiders that engage in non-routine trades that earn abnormal returns. Other proxies include insiders with short investment horizons (Akbas, Jiang, and Koch, 2020) or who trade profitably ahead of earnings announcements (Ali and Hirshleifer, 2017).[2] Our measure of the conditional probability of an insider's propensity to trade on private information is positively correlated with these measures, but the mixture model estimates contain a significant amount of independent information. In particular, controlling for whether an insider is a non-routine trader, a short-horizon trader, or makes more profitable trades ahead of earnings announcements does not impact the out-of-sample

[2] Other examples in this vein include Cline et al. (2017), Biggerstaff et al. (2020), and Goldie et al. (2023). Cline et al. (2017) also uses an insider's return history to classify persistently profitable insiders, but their classification does not utilize information from the cross-section of insiders nor the noise in an individual insider's trading history.

return predictability of the mixture model conditional probabilities. We show how our method can be generalized to incorporate information in existing proxies. These generalized models reveal that insiders not classified by these proxies are, on average, as informed as those the proxies categorize as informed.

Our ex-ante estimates of the likelihood that a given insider engages in informed trade allow for improved classification of whether any *single* trade is informed or not. In particular, we use a trade-level mixture model implied by the insider-level mixture model that utilizes an insider's return history in order to infer whether a trade is informed or not. The model results in an informed trade classification threshold that is customized based on the insider's return history. We are also able to use information from the full cross-section of insiders to classify trades by insiders with short or even no trading history.

The estimation yields several empirical findings about informed insider trade. First, the prevalence of likely-informed buys is about twice as high as likely-informed sales. Second, the return thresholds necessary in order to classify trades as likely informed are often quite high. This helps explain why the SEC pursues relatively few cases against corporate insiders despite empirical evidence that some insiders' trades predict future returns. Third, classification becomes more precise for insiders with longer trading histories. Through several case studies, we show how different return histories for insiders trading the same security at the same time result in differences in clearing a given (statistical) burden of proof.

A vast literature studies whether trades made by corporate insiders contain information relevant to future stock returns.[3] Within this literature, a number of papers document that some insiders are more likely to make informed trades than others (e.g., Cohen et al., 2012; Akbas et al., 2020). Our work relates to this literature but focuses on a different economic question. These papers establish that the trades of insiders that behave in ways the authors conjecture are related to opportunistic trading do, in fact, contain information on average. However, it is quite possible that other trading patterns would also identify informed trading. So, while these papers convincingly show that some trades have information, they are less able to speak to the universe of informed trading, either in terms of the fraction of insiders that take advantage of private information or in terms of the

[3]For example, research over the last fifty years includes Jaffe (1974), Seyhun (1986), Seyhun (1992b), Jeng, Metrick, and Zeckhauser (2003), Ravina and Sapienza (2010), and Cziraki and Gider (2021).

fraction of overall trades that are informed. A contribution of our paper is our ability to estimate a conditional probability that a given trade is informed for *all* trades disclosed by US corporate insiders in securities with publicly observable prices.

Our paper also contributes to the large literature on the optimal design and enforcement of insider trading regulations. For example, our results shed light on what fraction of trades are informed and what fraction are uninformed liquidity trades, potentially allowing regulators to evaluate the efficacy of disclosure rules in the context of Huddart et al. (2001), which provides theoretical predictions for the impact of insider disclosure on their trading behavior. DeMarzo et al. (1998) discuss optimal enforcement of insider trading and develop optimal investigation policies that depend on the number of shares bought/sold and the return to the stock, which are assumed to be the only observable information prior to a formal investigation. Our results suggest there is substantial information about the likelihood that an insider used private information in both the history (or lack thereof) of the insider's average trading performance and its noisiness, as well as the cross-section of performance across insiders and trades. This could impact the theoretically optimal enforcement policy, improving enforcement efficiency.

It is worth emphasizing a distinction between our study of the economic informativeness of disclosed trades by corporate insiders and the literature studying illegal insider trading (e.g., Meulbroek, 1992; Ahern, 2017; Kacperczyk and Pagnotta, 2019). Such studies primarily concern trades made following tipped information and do not necessarily directly involve corporate insiders. Trade by corporate insiders is illegal if it is based on material, nonpublic information (17 CFR 240.10b-5).[4] It is possible that some of the trading activity in our study is, in fact, illegal, but whether the economic materiality we document amounts to legal materiality is beyond the scope of this article.

Finally, another large literature discusses the costs and benefits of insider trading more generally. Going back to at least Hirshleifer (1971), many papers have identified tradeoffs that can arise from allowing insiders to trade on their information. For example, Dye (1984) points out that firm managers are often encouraged to buy stock in their firms and, therefore, suggests insider trading may allow for the efficient provision of incentives that outweigh adverse selection effects. Leland (1992) discusses how price efficiency resulting from privately informed insider trades may allow

[4]For an empirical analysis of the effectiveness of regulation on corporate insider-trading activity, see Seyhun (1992a).

for better resource allocation in the firm, with value improvements that again may outweigh the costs of insider trading. Theoretical work in this area is traditionally challenging to test because informed trading by insiders is largely unobservable. Our results take a step in this direction by providing a methodology for identifying and quantifying privately informed trades by corporate insiders.

2. Detecting Which Insiders Trade on Information

2.1. Modeling the Cross-section of Insiders as a Mixture Distribution

We model the distribution of average insider abnormal returns as a mixture of two distributions: an uninformed distribution and an informed distribution. A fraction $1 - \pi$ of insiders make trades that are uninformed. The remaining fraction π of insiders make trades that are, on average, informed. Empirically, the econometrician is able to estimate average abnormal returns for a given insider, denoted $\bar{r}_i$. The dispersion in estimated average abnormal return across insiders belonging to either group is driven by two components: true variation in informed trading and estimation error. Denote the true average abnormal return of insider i by α_i. We assume that the (unobservable) true average abnormal return of uninformed insiders is a point mass at zero ($\alpha = 0$) and that the true average abnormal return of informed insiders is distributed exponentially with mean μ ($\alpha_i \sim \text{Exp}(1/\mu)$).[5] The estimated abnormal return $\bar{r}$ is measured with estimation error, e_i, which is assumed independent of α_i and normally distributed around zero with a standard deviation of s_i, the standard error of insider i's abnormal performance. Thus, the estimated abnormal performance is $\bar{r}_i = \alpha_i + e_i$.

Figure 1 provides an illustration of the mixture model. Panel (a) shows the relative frequencies of the unobservable true abnormal return of insiders. Insiders that do not make informed trades comprise the grey bin located at zero, while the remaining π insiders make informed trades with magnitudes of varying amounts (the hatched purple bins). Thus, the unconditional distribution of informed insider trading is a mixture of the uninformed and informed component distributions.

[5] It is possible that only a strict subset of an informed insider's trades are informed. For instance, a possible specification of trade-level returns is $r_{ij} = y_{ij}\alpha_i + \varepsilon_{ij}$, where y_{ij} is Bernoulli with probability p_i, α_i follows some positive distribution f_α, and ε is normally-distributed. We take a quasi-maximum-likelihood approach to the problem. Our assumption of an exponentially distributed α for an informed insider's average return is an approximation of the complicated mixture distribution that would result from averaging a sample of trade returns with $p_i \in (0, 1)$. For instance, if α_i were exponentially distributed, then α_i would follow a zero-inflated hyper-exponential distribution.

Panel (b) of Figure 1 shows the effects of estimation noise on these component distributions. With noisy measures of the true informed insider trading, the distributions of $\bar{r}$ for uninformed and informed insiders overlap. The distribution of $\bar{r}$ for uninformed insiders is normally distributed around zero; all variation is due to estimation noise. The distribution of $\bar{r}$ for informed insiders is the convolution of an exponential and normal random variables;[6] variation in this distribution is due both to variation in the degree of informed trading and variation due to estimation error. In the example, the substantial overlap in the distributions leads to an unconditional distribution that is unimodal with positive skewness.

Let $f_{\mathrm{I}}(\bar{r}_i|\text{informed})$ denote the density of the observed average abnormal return conditional on an insider trading on information. Under the assumptions that the average profitability from informed trading is exponentially distributed and estimation noise is normally distributed, the conditional density of $\bar{r}$ is:

$$f_{\mathrm{I}}(\bar{r}_i|\text{informed}, s_i) = \int_{-\infty}^{\infty} g(\bar{r}_i - a; \mu) \cdot \phi(a; s_i) \ \mathrm{d}a \,, \tag{1}$$

where $\phi(\cdot; s_i)$ is the density of a mean-zero normal variable with standard deviation s_i and $g(\cdot; \mu)$ is the density of an exponential variable with mean μ. The unconditional density function for insider i's estimated average abnormal return $\bar{r}$ is:

$$f(\bar{r}_i|s_i) = (1 - \pi) \cdot \phi(\bar{r}_i; s_i) + \pi \cdot f_{\mathrm{I}}(\bar{r}_i|\text{informed}, s_i) \,. \tag{2}$$

The parameters of the model are π and μ. The likelihood function L for a sample of average abnormal returns of N insiders is:

$$L\left(\bar{r}_1, \bar{r}_2, ..., \bar{r}_N|s_1, s_2, ..., s_N, \pi, \mu\right) = \prod_{i=1}^{N} f(\bar{r}_i|s_i) \,. \tag{3}$$

To estimate the parameters π and μ, we maximize (3) subject to the restrictions that $\pi \in [0, 1]$ and $\mu > 0$.

[6] This random variable is known as an exponentially-modified gaussian random variable. Its density admits a closed-form expression, which reduces the computational burden of estimating the model.

2.2. Conditional Probabilities and Expectations

Given estimates for π and μ, the model allows calculations of the conditional probability that a particular insider i is informed, conditional on the insider's realized average abnormal return $\bar{r}_i$, its standard error s_i, and the estimated parameters. Denote the conditional probability by $\tilde{\pi}$. The conditional probability that insider i makes informed trades is:

$$\begin{aligned}\tilde{\pi}_i =& \Pr\left(\text{insider } i \text{ trades on information} | \bar{r}_i, s_i, \pi, \mu\right) \\ =& \frac{\pi \cdot f_{\mathrm{I}}(\bar{r}_i|\text{informed})}{(1-\pi)\cdot\phi(\bar{r}_i; s_i) + \pi \cdot f_{\mathrm{I}}(\bar{r}_i|\text{informed})}\,.\end{aligned} \tag{4}$$

Let $\tilde{\mu}_i$ denote the conditional expectation of insider i's information *conditional* on belonging to the informed component distribution (along with parameter values and realized $\bar{r}_i$ and s_i). Note that the conditional expectation of insider i's information is zero if i belongs to the no-informed-trading component. Thus, the conditional expectation of the magnitude an insider trades on information, conditional on their average abnormal return, standard error, π, and μ, which we denote $\tilde{\alpha}_i$, is:

$$\begin{aligned}\tilde{\alpha}_i =& \mathbb{E}\left[\alpha_i | \bar{r}_i, s_i, \pi, \mu\right] \\ =& \tilde{\pi}_i \tilde{\mu}_i\,.\end{aligned} \tag{5}$$

Let $f_{\alpha^+|\bar{r}}$ denote the true density of their abnormal return, α, given the insider uses information ($\alpha > 0$) and after observing the average return, $\bar{r}$, from their prior trades. The conditional expectation of insider i's average profitability, $\tilde{\mu}_i$, conditional on being in the component distribution that utilized private information, is calculated:

$$\begin{aligned}\tilde{\mu}_i =& \mathbb{E}\left[\alpha_i | \bar{r}_i, s_i, \pi, \mu, \text{informed}\right] \\ =& \int_{-\infty}^{\infty} a \cdot f_{\alpha^+|\bar{r}_i}(a|\bar{r}_i)\ \mathrm{d}a\,.\end{aligned} \tag{6}$$

We show in Appendix A that, under our distributional assumptions, $f_{\alpha^+|\bar{r}_i}$, is a normal distribution with mean of $\bar{r}_i - s_i^2/\mu$ and standard deviation of s_i that is truncated below at zero. Thus, $\tilde{\mu}_i$ is the mean of this truncated normal distribution.

Figure 2 illustrates the conditional probability (4) and conditional expectation (5) as a function of average abnormal return $\bar{r}_i$ and estimation noise s_i. Consistent with intuition, both are increasing functions of the average abnormal return.

The effect of estimation noise is more interesting. In Panel (a), see that the amount of estimation noise in the average abnormal return substantially affects inference about whether a particular insider trades on information. For low levels of estimation noise, the average abnormal return is a fair proxy for whether the insider trades on information. Negative abnormal returns are more likely to be uninformed insiders, while positive abnormal returns are more likely to be informed insiders. As estimation noise increases, however, the average abnormal return is a less reliable proxy for whether the insider trades on information. The slope of the conditional probability function is much shallower in average abnormal return, consistent with the fact that the realized average return could be high or low due to estimation error (i.e. luck) rather than true trading on information.

The effects of estimation noise for the conditional expectation are also interesting. The conditional expectation is a convex function of realized average abnormal returns, with greater convexity for insiders with less estimation noise. For low noise, the conditional expectation is not far from simply taking the maximum of $\bar{r}_i$ and zero. The shape of the conditional expectation function is flatter with greater estimation noise. This is because some insiders that truly trade on information may have been unlucky and realized a negative $\bar{r}_i$. Similarly, some insiders that do not trade on information may have been lucky and realized a positive $\bar{r}_i$. The mixture model approach essentially shrinks these realized returns as a function of the estimation noise.

2.3. Data

The data on stock transactions by corporate insiders is from the Thomson Reuters Insider Filing database, which captures and cleans Form 4 filings by corporate insiders. Our sample covers trades from 1985 to 2022. We also use stock returns and trading volumes from CRSP and financial reporting information from Compustat.

On a given transaction date, insiders sometimes report multiple transactions in a single stock and/or across multiple stocks. We aggregate such trades to the daily level to create an insider-stock-date panel. Index insider i's trades by $j = 1, ..., n_i$.[7] For trade j made by insider i on day t,

[7] In our full-sample estimation, n_i is simply the total number of distinct stock-date observations for insider i. In our out-of-sample estimation, we estimate an annual time-series of π and μ using only past available data. For this analysis, n_i is the total number of distinct stock-date observations for insider i as of the year end of the estimation.

we calculate a 21 trading day market-adjusted abnormal return

$$r_{ij} = D_{ij} \cdot \left(\prod_{k=1}^{21} (1 + r_{j,t+k}) - \prod_{k=1}^{21} (1 + r_{m,t+k}) \right) , \tag{7}$$

where $r_{j,t+k}$ is the day $t+k$ return of the stock purchased or sold in trade j, $r_{m,t+k}$ is the day $t+k$ CRSP-value-weighted return, and D_{ij} denotes a buy sell indicator defined as:

$$D_{ij} = \begin{cases} +1 & \text{for purchases} \\ -1 & \text{for sales} . \end{cases} \tag{8}$$

The mixture model described in Section 2.1 uses an average abnormal return and its standard error for each insider i as inputs to estimating parameters π and μ. We calculate an average abnormal return for insider i as:

$$\bar{r}_i = \frac{1}{n_i} \sum_{j=1}^{n_i} r_{ij} , \tag{9}$$

as well as its associated standard error s_i. To ensure sufficient data for estimation, we require an insider have at least 10 stock-day observations to be included in the mixture model estimation. To limit the effect of outliers, the sample is trimmed at the 1% and 99% levels of average abnormal returns.

Figure 3 plots histograms of the average abnormal returns, $\bar{r}_i$, in Panel (a) and of the standard errors s_i in Panel (b), and Table 1 reports distributional statistics. While the mode of the $\bar{r}$ distribution is close to zero, the cross-sectional average $\bar{r}$ is positive 66 basis points. Just over half (55%) of the insiders have positive average abnormal returns. The $\bar{r}$ distribution exhibits slight positive skewness consistent with some insiders trading on information. Panel (b) of Figure 3 and the second column of Table 1 show there is substantial variation in the amount of estimation noise in $\bar{r}$. The cross-sectional average standard error is 2.5% and the cross-sectional standard deviation is 1.75%. This suggests value in using an informed insider classification designed to explicitly account for estimation noise like the mixture model described in Section 2.1.

2.4. Insider-Level Empirical Prevalence of Informed Insider Trading

Table 2 reports estimates of the mixture model described in Section 2.1 on the trimmed sample of insiders with at least 10 trades. The empirical estimate of the fraction of insiders trading on information, π, is 28.6%. The average magnitude of information for this set of insiders, μ, is 3.6%.

How can these estimates be used to consider whether a particular insider uses information? Given a particular insider's average past abnormal return and its standard error, we can calculate the conditional probability $\tilde{\pi}$ (4) and conditional expectation $\tilde{\alpha}$ (5) using the full-sample estimates of π and μ. These estimates take into account both the noise inherent in the particular insider's past history of abnormal returns and information from the cross-section of insider trades used to estimate π and μ.

Figure 4 displays how the conditional probability and expectation vary as a function of an insider's estimated abnormal return and its noise (i.e., its standard error). To see how the mixture methodology takes estimation noise into account, consider an insider whose past trades have averaged an abnormal return of 2.5%, which is close to the 75th percentile of average returns. Panel (a) of Figure 4 shows that an econometrician (or regulator) should make very different inference about the likelihood that this individual uses private information depending on the standard error of that 2.5% average abnormal return. If that insider's noise is in the top quartile, it is more likely that the insider *does not trade on private information* even with such a high return. It is only for noise levels below the first quartile that the econometrician should think it more likely than not that the insider with average returns in the top quartile trades on private information.

Put differently, consider the average abnormal return a regulator would need to observe in order to consider an insider's conditional probability of trading on information to be 80%. For low noise, say at the 5th percentile of the empirical distribution, an abnormal return of about 2% would lead to an 80% conditional probability. The required average abnormal return jumps to over 3% for the first quartile of noise and to 5% for median noise. If the average abnormal return is very noisy at say, the 75th percentile, the insider would need to average an 8% abnormal return on their trades to reach an 80% conditional probability.

Panel (b) of Figure 4 illustrates the associated effect of the noise-reduction in estimating the conditional expectation of the average informativeness of an insider's trades. In order to conclude that an insider exhibits information of 2% on average, one would need to observe an empirical average abnormal return for the insider of about 3% for an insider with 1st quartile noise; to reach the same conclusion for 3rd quartile noise, an insider would need to have earned an abnormal return of over 5%. As the noisiness of the estimated average abnormal return rises, the model shrinks the conditional expectation more because the conditional probability that the insider uses information is lower.

In addition to the full-sample estimation, we also estimate the model on expanding windows. Specifically, the mixture model is estimated each year using the latest average abnormal return and standard error for each insider with at least ten trades prior that year end. The time-series of π and μ are plotted in Figure 5. The parameter estimates are fairly stable over time. Aside from the first few years of the sample, the fraction of insiders utilizing information has consistently been around the 30% range. The average magnitude of the utilized information is more variable, dropping from above 5% in the late 1980's to below 4% in the mid-1990s before rising to almost 5% again in the early 2000s. After that, there has been a fairly steady decline in the magnitude of utilized information.

2.5. Characteristics of Informed Insiders

Insiders are not homogeneous. There is the potential for insiders to differ in terms of incentives, ability, or even information sets. For example, the CEO likely has a different information set than the board chair, despite both clearly having access to important private information about the firm. It is therefore natural to ask whether the prevalence of informed trading by insiders varies as a function of the insider's role in the firm. Using the roles in the firm as disclosed on Form 4s, we test whether the insider's conditional probability of being informed ($\tilde{\pi}$) and their conditional expected alpha ($\tilde{\alpha}$), as estimated based on insider average returns as in Section 2 and expanding window estimates of π and μ, are a function of their role.

To do this, we regress $\tilde{\pi}$ and $\tilde{\alpha}$ on indicators for each of the 54 possible roles indicated in the Thomson Reuters Insider Filing database. Thomson reports up to four roles that can be indicated on a given filing. It is therefore possible that for any given filing, more than one indicator is turned on. An indicator is set to one if the insider reported that role on any filing in a calendar year, and zero otherwise.

We conduct two analyses and report the results in Table 3. The first aggregates these dummy variables into broader groups: C-suite, directors, owners, non-C-suite officers, non-officer managers, and others. The results from this analysis are reported in Columns (1) and (3) of Table 3. The second analysis disaggregates these into indicator variables for each of the individual roles. We report these results in Columns (2) and (4). For brevity, we only report results for the roles that appear in at least 5% of the observations. All other role indicators are included as controls, but the coefficients are suppressed for space.

C-suite executives are more likely to make informed trades and their expected alpha is substantially higher. Specifically, insiders from the C-suite are 3 percentage points more likely than than those not in the C-suite to be drawn from the distribution of investors trading on private information. This is roughly 10% greater than the unconditional π estimate of approximately 30%. These investors have an expected monthly alpha ($\tilde{\pi}$) roughly 23 bps larger than insiders not in the C-suite. Inside owners have the largest probability of making informed trades and the largest expected performance (7% and 83 bps, respectively), while non-C-suite executives are actually less likely to make informed trades.

When we disaggregate these roles, it is clear that the C-suite result is driven almost entirely by CFOs (Wang et al., 2012) and the ownership result stems from large insider blockholders. On the other hand, the lower probabilities and performance expected from non-C-suite officers is prevalent across most of those roles. Finally, while the aggregate director category is insignificant, this includes directors that may take on other roles as well. When we look at directors that serve no other roles, the results suggest a lower propensity to engage in informed trades.

3. Out-of-Sample Predictability and Learning by Market Participants

In this section, we consider the out-of-sample performance of our mixture model estimates. To do so, we use the annual π and μ time series estimated using expanding windows described above to calculate a conditional expectation (5) of insider informed trade for each insider with at least 10 distinct trading days prior to a given year end. Insiders are sorted into quintiles on the basis of the conditional expectation each year. We then test whether trades made by insiders with higher conditional expectations predict returns using both predictive regressions and portfolio analysis.

3.1. Regression Analysis

To test the model out-of-sample, we consider whether buys and sells by insiders with different lagged conditional expectations predict future stock performance. Specifically, we create a stock-month panel with indicator variables for whether there were any purchases or sales by an insider classified in a particular quintile as of the prior year-end. For instance, Buy Quintile 4 (Sell Quintile 4) is an indicator variable for whether any insider in the top quintile of conditional expectation bought (sold) shares in month t. We regress month $t+1$ stock returns on buy and sell indicator variables for each quintile of conditional expectation. Note that this is an out-of-sample exercise of our ability to rank insiders' propensity to use information because the quintile is formed using information known as of the beginning of month t.

Table 4 reports the estimates of the regression of future monthly returns on buy and sell indicators for each quintile of insider conditional expectation. As is standard, we control for a stock's market capitalization, book-to-market ratio, and lagged monthly and annual return. We consider specifications both with (even numbered columns) and without (odd numbered columns) month fixed effects

There is substantial cross-sectional spread in future returns as a function of buying activity by insiders across conditional expectation quintiles. Without month fixed effects, the difference in future returns for stocks with buying activity by top-quintile insiders relative to bottom-quintile insiders is 82 basis points per month, or almost 10% per year. The predictability remains strong with the inclusion of month fixed effects. The Hi-Lo spread is 69 bps per month, or 8.3% annually. The differences are statistically significant at the 1% level.

Selling activity also results in cross-sectional spread in future returns as a function of an insiders ex-ante conditional expectation quintile (columns (3) and (4)). The Lo-Hi spread is about 46 bps per month without monthly fixed effects and 31 bps per month with fixed effects. The differences are again statistically significant.

The spreads in future stock returns resulting from both buying and selling activities remain practically unchanged and strongly significant if we include buying and selling indicators in the same regression (columns (5) and (6)). Overall, the results of Table 4 provide strong support for the mixture model's ability to differentiate between insiders with higher propensities to profit from their private information.

3.2. Portfolio Analysis

An alternate way to test the predictive power of the mixture model estimates is assess the abnormal performance of portfolios formed based on the trading activity of insiders with different ex ante conditional expectations. We form portfolios by double-sorting insider-stock pairs into (1) quintiles of an insider's lagged conditional expectation and (2) quintiles of an insider's signed trading activity. The signed trading activity for an insider-stock month is their signed order flow (i.e., the sum of all the insider's purchases minus the sum of their sales) scaled by the overall trading volume in the stock that month. Each month, we form 25 portfolios based on the intersection of the two sets of quintiles. The portfolios are equal-weighted combinations of the insider-stock pairs falling in a given conditional expectation quintile-insider order flow quintile sort. The portfolio is rebalanced each month based on the trading activity of insiders in that month. Note this is an out-of-sample exercise as the conditional expectation quintiles are formed using information about the insiders that is available as of the prior year end for a given portfolio formation month.

Table 5 reports average excess returns (Panel A) and abnormal returns under various benchmark models (Panels B, C, and D) for each of the 25 portfolios. The benchmark models are a market model benchmark, a three-factor Fama and French (1993) model augmented with a momentum factor (Carhart, 1997), and the five-factor Fama and French (2015) plus momentum model. For each conditional expectation quintile, the table also reports returns of the High Minus Low Insider Order Flow portfolio.

Across all return measures and all conditional expectation quintiles, the point estimates of the High Minus Low Insider Order Flow portfolio are positive, consistent with stronger buying activity by insiders predicting better subsequent performance than their selling activity. For the lowest quintile of insider conditional expectation, the economic magnitude is small and statistically indistinguishable from zero. The economic magnitude of the High Minus Low abnormal performance is monotonically increasing in the quintile of conditional expectations. Regardless of the benchmark model used, the economic magnitude of the High Minus Low abnormal performance is about 1.5% per month for the top quintile of conditional expectation. These results show that the mixture model is able to separate insiders that are more or less likely to trade using their information advantage.

3.3. Learning by the Market

We have shown that an econometrician can learn through a trader's return history, in conjunction with information from the cross-section of insider histories, which insiders are more or less likely to trade on information. Indeed, other proxies for this heterogeneity exist (e.g., whether an insider routinely trades in the same month each year as in Cohen et al. (2012)). We now turn to the extent to which the market has learned about this heterogeneity over time.

Figure 6 shows cumulative returns from the monthly portfolio strategies described in the previous section. The top panel reports cumulative returns for hedge portfolios that buy stocks with strong inside buying pressure and sell stocks with strong inside selling pressure. The black solid (red dashed) line represents this strategy for insiders in the top (bottom) quintile of ex ante conditional expectation. The bottom panel reports the cumulative performance for hedge portfolios that either (1) buy the top $\tilde{\alpha}$ quintile's strong buys and shorts the bottom $\tilde{\alpha}$ quintile's strong buys (black solid line) or (2) buy the top $\tilde{\alpha}$ quintile's strong sells and shorts the bottom $\tilde{\alpha}$ quintile's strong sells (red dashed line), or (3) buys the first hedge portfolio of strong buys and shorts the second hedge portfolio of strong sells (blue dashed-dotted line).

Visual inspection of the cumulative returns indicate that these portfolio strategies performed well for the first 20 years or so of the sample, but that the performance has been flatter since around 2012. Interestingly, 2012 was the publication year of the influential Cohen et al. (2012) paper. Indeed, in untabulated results, we confirm that there is no statistically significant performance differentials at the monthly frequency post-2012.

This raises the question of whether the prevalence of informed insider trades has declined over time. To answer this, we estimate the insider-level mixture model on 10-year rolling windows. That is, we take insider-level average abnormal returns using the trades in a given 10-year window. The resulting estimates are plotted in the Internet Appendix. The fraction of insiders using information was fairly steady around 30% for the first 20 years of the sample before declining to around 25%. The mean of informed trading has varied more over time, but it has also been lower over the past decade. However, there is still a substantial amount of informed trade by corporate insiders. Over the decade ending in 2022, fully 25% of insiders are estimated to make informed trades.

What then can explain the lack of predictability at the monthly frequency over the last decade? We hypothesize that increased information acquisition and learning by market participants about which insiders trade on information has reduced the time horizon for market prices to reflect information embedded in insider trading activity. Below we test this hypothesis. First, we show that portfolios double-sorted on an insider's ex ante conditional expectation and buying or selling indicators do remain profitable post-2012, but only if the portfolios are implemented much more frequently than the monthly horizon typically employed in the insider trading literature. Second, we provide evidence of dramatically increased information acquisition by sophisticated financial institutions around the same time the monthly strategy performance declines.

3.3.1. Daily Portfolios

Figure 7 reports cumulative returns of portfolios analogous to the monthly portfolios considered in Figure 6 but with portfolio formation occurring the day after the trade date.[8] A stock remains in a portfolio for forty days. The other difference from the analysis in Figure 6 is that the portfolios in Figure 7 do not condition on monthly volume; the second sort is simply whether the trade is a buy or a sell.

Unlike the monthly strategies, daily portfolios formed the day after an insider's trade continue to exhibit strong performance even post 2012. This is confirmed statistically in Table 6 which reports risk-adjusted performance. The reported returns in the table are expressed as monthly returns for comparability to the monthly analysis. In the post 2012 sample, insider buys have outperformed insider sells in all quintiles of ex ante conditional expectation, and the differential is larger for the higher conditional expectation quintiles. Buys made by the highest $\tilde{\alpha}$ quintile outperform buys of the lowest $\tilde{\alpha}$ quintile while high $\tilde{\alpha}$ quintile sells underperform low $\tilde{\alpha}$ quintile sells. The ex ante classification of insider informed trade continues to have out-of-sample validity even in the most recent decade.

The discrepancy between the monthly and daily portfolio profitability over the last decade suggests that the information embedded in trading by corporate insiders is incorporated quickly into prices. To gauge how quickly, we vary the delay with which a stock enters the daily portfolio. Specifically, we consider portfolios that wait 1, 3, or 5 days post-trade to include a stock in the portfolio. In each case, the stock exits the portfolio forty days after the insider's trading date.

The cumulative returns of these three portfolio formation delays are shown in Figure 8. For both the high $\tilde{\alpha}$ quintile buy minus sell strategy (top panel) or the high $\tilde{\alpha}$ buys minus the low $\tilde{\alpha}$ buys (bottom panel), the profitability is lower if more time elapses before a stock enters the portfolio.

[8] Note that trades need not be disclosed until 2 days after trade, so this portfolio is not tradeable. Our objective is to show that the information embedded in insider trades is still economically material, not to demonstrate whether an investor could implement this trade.

3.3.2. Information Acquisition

In this section, we provide evidence of dramatic increases in information acquisition of insider trade disclosures by sophisticated investors. Several recent papers study information acquisition of public disclosures from the SEC website and their relation to investment performance and market efficiency. Chen, Cohen, Gurun, Lou, and Malloy (2020a) show that institutions track insider trading filings of the same firm and purchases of tracked firms outperform non-tracked purchases. Chen, Kelly, and Wu (2020b) show that hedge funds increase information acquisition following declines in analyst coverage, mitigating reductions in market efficiency. Crane et al. (2023) show that hedge fund performance is higher for funds that access more public filings. We follow the methodology in Crane et al. (2023) in identifying institutional information acquisition of Form 4 insider trading disclosures from the SEC EDGAR search logs.

Given the documented change in the time horizon at which insider trades predict returns around 2012 (the publication date of Cohen et al. (2012), we are interested in changes in information acquisition around this time. We focus on financial institutions that engage in large amounts of EDGAR activity (defined as having at least 200 days of EDGAR activity within at least one year and at least 50 daily downloads each day within that year). We identify a number of firms that exhibit dramatic increases in direct information acquisition of insider trade disclosures from EDGAR in the time period preceding publication of Cohen et al. (2012).

Figure 9 shows time-series of the number of weekly downloads of insider trade reports for a set of financial institutions that exhibited large increases leading up to 2012. These firms include some of the most sophisticated asset managers, including firms like D.E. Shaw and Renaissance Technologies. The structural break for some of these firms entails an initial set of extremely large-scale downloads (likely for historical database creation) followed by a fairly steady increased level of periodic downloads. Given this increased information acquisition by sophisticated investors, it is not surprising that the information embedded in insider trade disclosures is more quickly incorporated into prices over the subsequent decade.

4. Relation to Existing Proxies of Informed Corporate Insiders

As discussed in the introduction, the existing literature has produced proxies for which insiders are more or less likely to trade on private information. In this section, we compare our conditional expectation measure to several of the most prominent of these proxies: non-routine traders (Cohen, Malloy, and Pomorski, 2012), short horizon insiders (Akbas, Jiang, and Koch, 2020), and profitability of trades ahead of quarterly earnings announcements (Ali and Hirshleifer, 2017). Non-routine traders are insiders who have made at least one trade in each of the past three years, but who do not have trades in a particular calendar month in each of the three years. Short horizon insiders are those whose trade direction is not consistently in the same direction over the prior ten years. An insider who always buys or always sells is classified as a long-horizon insider, Insider with some buying and selling activity within the year are classified as either medium- or short-term insiders. High QEA profitability insiders are those that trade ahead of quarterly earnings announcements and whose pre-QEA trades are associated with the highest quintile of QEA-window profitability.

We first show how the mixture model estimates relate to existing measures and that our estimates provide incremental information about the informativeness of insider informed trading. We then show how our methodology can be generalized to incorporate these alternative proxies.

4.1. Comparison to Existing Proxies of Informed Corporate Insiders

We first consider how the mixture model conditional probability (4) estimates correlate with these measures. The results are reported in Panel A of Table 7. Consistent with the conclusions of the prior literature, we find positive correlations between each proxy and the conditional probability $\tilde{\pi}$. Non-routine insiders have 4% higher $\tilde{\pi}$, compared to an average of 22% for routine insiders. Long-horizon insiders have a 22% conditional probability. With increasing levels of short-termism, the conditional probability rises, consistent with Akbas et al. (2020). Medium horizon insiders have 4% higher $\tilde{\pi}$, and short-horizon insiders have 9% higher $\tilde{\pi}$. Finally, insiders in the top quintile of QEA profitability exhibit $\tilde{\pi}$s that are on average 5% higher than the remaining four quintiles.

Each of the proxies is also associated with higher mixture model conditional expectations ($\tilde{\alpha}$ s) as well (Panel B of Table 7). The average conditional expectation of non-routine insiders is about 50% higher than that of routine insiders. In terms of abnormal returns, non-routine insiders' average $\tilde{\alpha}$ is 29 bps higher than the 60 bps average $\tilde{\alpha}$ of routine insiders. The wedge is even larger for long- versus short-horizon insiders. Long-horizon insiders have average $\tilde{\alpha}$ of 57 bps; short-horizon insiders' average $\tilde{\alpha}$ is twice as much, with medium-horizon insiders falling in between. Top quintile QEA profitability insiders exhibit conditional expectations that are 33 bps higher than those of the remaining four quintiles.

It is important to note that, while these proxies are positively related to our insider-level measure of informed trading as expected, they actually explain very little of the cross-sectional variation in informed trading as measured by $\tilde{\pi}$ and $\tilde{\alpha}$. This suggests that our measure is capturing different information than prior work.

4.2. Predictability Controlling for Existing Measures

Given the positive correlations with existing proxies of informed trade and the conditional probabilities, it is important to ask whether the mixture model provides incremental information about the informativeness of insider informed trading. To assess this, we add indicators to our stock-month return panel for buying and selling activity by insiders classified as informed by the alternative approaches. The results are presented in Table 8. We include time fixed effects and standard control variables, but suppress their coefficient estimates for space considerations.

Columns (1), (3), and (5) confirm the results in the existing literature that trading activity by insiders classified as informed by existing measures predicts future stock returns. Consistent with Cohen et al. (2012), stock returns are about 50 bps higher in months following a non-routine purchase and 17 bps lower following a non-routine sell, while buys and sells by routine insiders do not predict returns. Consistent with Akbas et al. (2020), stock returns are higher (lower) following purchases (sales) by shorter horizon investors than long-horizon investors. Purchases by short-horizon investors result in an 83 bps higher stock return in the next month than purchases by long-horizon investors. Similarly, sales by short-horizon investors result in a 36 bps lower stock return in the next month than sales by long-horizon investors. Finally, purchases by insiders in the top quintile of QEA profitability are 66 bps higher than those of the remaining quintiles.

In Columns (2), (4), and (6) of Table 8, we assess whether the predictability of the mixture model classification documented in Table 4 controlling for these measures. The cross-sectional differences in future monthly returns following purchases or sales by high and low conditional expectation quintiles remains statistically and economically significant. For purchases, the Hi-Lo spread ranges between 62 and 69 bps per month, or 7.4 to 8.3% annually. For sales, the Lo-Hi spread ranges from 26 to 30 bps per month, or 3.1 to 3.6% annually. These differences are of a similar magnitude to those reported in the last column of Table 4, indicating that controlling for existing proxies does not much affect the predictability of the mixture model. On the other hand, inclusion of trading indicators for the mixture model estimates substantially reduces the predictability of non-routine buys and sells as well as the high QEA profitability quintile trading indicators. The horizon measure of Akbas et al. (2020) continues to provide incremental information to the mixture model estimates, particularly for purchases.

Overall, the results of Tables 4 and 8 show that the ex ante conditional expectation quintiles for insiders from the mixture model contain economically and statistically significant information about which insiders are more likely to trade on private information.

4.3. Incorporating Existing Measures into the Mixture Model

It is possible to incorporate an alternate proxy into the mixture framework. An existing proxy either classifies an insider as (1) one who uses informed or (2) one who does not use information, or (3) one that is not classified either way due to the insider not satisfying sample screens. Denote indicator variables for these three mutually exclusive classifications as $\mathbf{1}_{\text{Informed},i}$, $\mathbf{1}_{\text{Uninformed},i}$, and $\mathbf{1}_{\text{Unclassified},i}$, respectively. In the mixture model, we parameterize the probability that an insider i is informed using these variables as:

$$\pi_i = \exp\left(\frac{\beta_1 \mathbf{1}_{\text{Informed},i} + \beta_2 \mathbf{1}_{\text{Uninformed},i} + \beta_3 \mathbf{1}_{\text{Unclassified},i}}{1 + \beta_1 \mathbf{1}_{\text{Informed},i} + \beta_2 \mathbf{1}_{\text{Uninformed},i} + \beta_3 \mathbf{1}_{\text{Unclassified},i}}\right). \tag{10}$$

This essentially results in distinct mixing probabilities for each classification.

Table 9 reports estimates of this model for each of the alternative proxies. Consistent with the existing literature, the estimated π_i is higher for the informed classification than the estimated π_i for the uninformed classification. The estimation provides two more novel facts. First, insiders that are unclassified based on existing work exhibit similarly high likelihoods of trading on information as those classified as informed. Second, while the non-opportunistic insiders have lower propensities to trade on information, their estimated π_i's are non-zero. A significant fraction of these insiders also engage in informed trade.

We use these proxies to parameterize π to demonstrate this generalization of the model and to compare to existing work. A regulatory agency may of course be interested in including additional characteristics of insiders, such as those in Section 2.5.

5. Detecting Which Trades Are Informed

5.1. A Trade-Level Mixture Model

In this section, we are interested in classifying whether an individual trade made by insider i was potentially privately informed or not. We will use the expanding window estimates of the insider-level mixture distributions π and μ. We will also utilize insider i's average return $\bar{r}_i$ and standard error s_i resulting from the insider's $j-1$ previous trades. We are interested in a probabilistic model of return r_{ij} with:

$$r_{ij} = \alpha_{ij} + \varepsilon_{ij}\,. \tag{11}$$

We assume that ε_{ij} is normally distributed with mean zero and standard deviation σ_i. For insiders with a reasonably large history ($j > 10$), we use the standard error of insider i's return history and set the trade-level standard deviation σ_i equal to $\sqrt{(j-1)} \cdot s_i$. For insiders with a relatively sparse history of trades ($3 \leq j \leq 10$), we set σ_i equal to a weighted average of the insider's historical σ_i and the cross-sectional median standard deviation, $\text{med}(\sigma)$, where the median is taken across insiders with at least 10 trades and the weight on the insider's standard deviation is $(j-2)/9$. For insiders making their first or second trade, we set σ_i equal to $\text{med}(\sigma)$.

If the insider has a past history of trades, we use it to inform the probability that α_{ij} is positive. Denote the probability that $\alpha_{ij} > 0$ as p_{ij}. For insiders trading for the first time, we estimate this probability using π, which implicitly utilizes information from the history of the cross-section of past traders. For insiders with any past history of trades, we use it in conjunction with π and μ to inform the probability that α_{ij} is positive. Specifically, the trade-level probability that $\alpha_{ij} > 0$ is:

$$p_{ij} = \begin{cases} \pi & \text{if } j = 1 \\ \tilde{\pi}_i & \text{otherwise}\,. \end{cases} \tag{12}$$

For this calculation, π and μ are estimated in expanding annual windows; we use trade-by-trade updating of $\bar{r}_i$ and s_i in calculating $\tilde{\pi}_i$ in equation (4).

The insider-level model informs the distribution of α_{ij} conditional on a positive α_{ij} realization. For insiders with a past history of trades, the conditional distribution of α_{ij} is $f_{\alpha^+|\bar{r}_i}$, which is a truncated normal distribution with normal mean of $\bar{r}_i - \sigma_i^2/((j-1)\cdot\mu)$ and standard deviation $\sigma_i/\sqrt{j-1}$. For insiders trading for the first time, the conditional distribution is simply the unconditional distribution of informed trade; that is, exponential with mean μ. The density of r_{ij} conditional on a positive α_{ij} realization is the convolution of α_{ij} and ε_{ij}:

$$h_{\mathrm{I}}(r_{ij}|\alpha_{ij}>0)=\begin{cases}\displaystyle\int_{-\infty}^{\infty} g\left(r_{ij}-a;\mu\right)\cdot\phi(a;\sigma_i)\;\mathrm{d}a & \text{if } j=1\\ \displaystyle\int_{-\infty}^{\infty} f_{\alpha^+|\bar{r}_i}\left(r_{ij}-a\right)\cdot\phi(a;\sigma_i)\;\mathrm{d}a & \text{otherwise}\,.\end{cases}\tag{13}$$

In Appendix B, we show that the convolution of a truncated normal and normal distribution for $h_{\mathrm{I}}(r_{ij}|\alpha_{ij}>0)$ with $j>1$ can be expressed in closed form, which substantially eases the computational burden of this trade-level model. The likelihood of observing return r_{ij} is a mixture:

$$h\left(r_{ij}\right)=(1-p_{ij})\cdot\phi(r_{ij};\sigma_i)+p_{ij}\cdot h_{\mathrm{I}}(r_{ij}|\alpha_{ij}>0)\,.\tag{14}$$

To classify whether a given trade was likely informed, we calculate the conditional probability of a positive α_{ij}:

$$\tau(r_{ij})=\frac{p_{ij}\cdot h_{\mathrm{I}}(r_{ij}|\alpha_{ij}>0)}{(1-p_{ij})\cdot\phi(r_{ij};\sigma_i)+p_{ij}\cdot h_{\mathrm{I}}(r_{ij}|\alpha_{ij}>0)}\,.\tag{15}$$

An econometrician (or regulator) can choose a threshold probability above which one classifies trade r_{ij} as potentially informed. For instance, the model suggests that trades with $\tau_{ij}>0.5$ are more likely than not to have been informed. Regulators with limited investigative and enforcement budget might choose a higher threshold, say 0.95, in considering which trades to investigate further. That is, regulators might investigate trades where the random variable $1(\tau_{ij}>0.95)$ equals one.

5.2. Conditional Probabilities and Return Thresholds

An important feature of this model for insider trades is that the classification threshold is customized to each insider based on their history of returns (as well as indirectly on the history of other insiders through the estimates of π and μ). We demonstrate graphically how the function τ varies as a function of an insider's historical average realized return, standard error, and number of trades.

Figure 10 plots the conditional probability (15) as a function of the realized trade return r_{ij}. The figure reports the probability conditional on an insider's past average abnormal return $\bar{r}_i$, the standard deviation of their past trades σ_i, and the number of past trades. As is natural, higher realized returns translate into higher conditional probabilities that a given trade was informed.

Each panel fixes the number of prior trades and their return standard deviation. Consider the top left panel, which considers an insider with a history of 10 trades with a standard deviation of 4%. An insider's past average return is informative about whether a given trade was informed. For an insider with a past average return of zero, the current trade return would need to be over 20% for the model to classify the trade as more-likely-than-not informed. This threshold is (lower) higher for insiders who have (gained) lost 1% on average on their past 10 trades.

As the number of prior trades increases, these differences increase and even higher returns are needed in order to classify the trade as informed (compare the left and right columns). This is because the conditional distribution $f_{\alpha^+|\bar{r}_i}$ is more tightly centered around $\bar{r}_i$ as the length of the history increases (holding σ_i constant).

The effect of an increase in the dispersion of an insider's past returns, σ_i, is to shift the conditional probability curves to the right (compare the top row to the bottom row of Figure 10). With more noise in the trader's history, the model requires a higher current trade return r_{ij} to reach the same conditional probability level that the current trade was informed.

An alternative way to visualize these relations is shown in Figure 11, which plots the trade-level return thresholds at which the current trade is classified as more-likely-than-not informed. This corresponds to a preponderance-of-evidence burden of proof. The top two panels plot these thresholds as a function of an insider's historical average return ($\bar{r}_i$) and the bottom two panels plot them as a function of the standard deviation of the insider's past trade returns, σ_i. As described above, the return threshold is declining in past average return. This decline is greater when there is a longer return history (compare the top left and top right panels of Figure 11). The return threshold is increasing in the standard deviation of past returns.

One thing to note from these figures is that for some histories, the return threshold can be fairly high. Insiders generally have a limited trading history and it is usually fairly noisy as well. For an insider with a past average return of zero with 10 trades that had a return standard deviation of zero, the current trade would need an abnormal return of almost 50% to be classified as more-likely-than-not informed. Of course, with a higher confidence level (say 95% rather than 50%), even more extreme trade-level returns are needed. Moreover, theory by Huddart et al. (2001) shows insiders may dissimulate, i.e. sometimes trade contrary to their information, if trades are disclosed publicly. Our results show that dissimulation will result in higher empirical return thresholds needed to classify a given trade as informed if dissimulation results in greater σ_i.

5.3. Trade-Level Conditional Probabilities of Informed Insider Trade

We calculate the trade-level conditional probability τ from equation (15) for a sample of *all* insider trades from 1991 through 2022. The probability depends on the population parameters π and μ estimated using expanding windows using all data up until the prior year end. These estimates are plotted in Figure 5.

Figure 12 plots histograms of insider-specific inputs to the conditional probability calculation: the trade abnormal return r_{ij} (Panel (a)), an insider's past average abnormal return $\bar{r}_i$ (Panel (b)), the standard deviation of an insider's past trades σ_i (Panel (c)), and the number of trades made by the insider prior to the given trade (Panel (d)). There is considerable heterogeneity in trade abnormal returns. There is a slight asymmetry toward positive abnormal returns, but the most populated bins are around zero abnormal return. The past average abnormal return distribution is naturally much less dispersed than the trade-level return distribution. In general, there is substantial dispersion in past trade abnormal returns with the modal σ_i being around 10%. The trade-level mixture model will take this underlying noise into account in estimating informed trade at the individual trade level. Finally, many trades must be classified without extremely long histories, which underscores the importance of using cross-sectional information embedded in π and μ in assessing individual trades.

Table 10 reports averages of the conditional probabilities τ and indicators for whether a trade is classified as very likely informed or uninformed as well as more-likely-than-not informed. The average conditional probability is 30%, which is close to the unconditional estimate of π reported in Table 2. If one takes a 5% critical value for classifying a trade as informed or uninformed (columns 2 and 4), then the model classifies 6.7% of trades as informed and 8.8% as uninformed. Unconditionally, 16.6% of trades are more likely than not to be informed. Figure 13 plots the density histogram of the trade-level conditional probabilities.

Consistent with prior work, the probability that a given trade is informed is strongly related to whether the trade is a purchase or a sale. The average conditional probability is higher for buys (35%) than for sales (27%), and over 10% of buys clear the 0.95 threshold compared to only 5% of sales. More sales are classified as likely-uninformed than buys (9.6% versus 7.2%).

Theory (DeMarzo et al., 1998) suggests that an optimal enforcement regime should condition on large trading volumes (or large price movements or both). In the spirit of their theory, we sort trades each month into quintiles based on a trade's signed trading strength, defined as the signed trade volume divided by overall trading volume in the stock that month. Consistent with theoretical intuition, the average conditional probability is higher for extreme buying and selling pressure. The frequency of likely-informed trades is J-shaped in signed trading strength, consistent with a higher prevalence of informed buying activity than selling activity.

Inferences about whether an insider's trades are informed or not get sharper with longer histories (i.e., greater numbers of past trades). Table 10 reports averages for four groups of the number of previous trades. For the groups with less than 50 trades, the average conditional probabilities are fairly similar, while the small set of insiders exhibiting more than 50 trades have a bit higher average τ. As the number of trades gets larger, however, the model classifies larger fractions of trades as likely-informed, likely-uninformed, or more-likely-than-not informed.

Figure 14 demonstrates how the inferences sharpen with trade history length in the data. In Panel (a), we see that the interquartile range of the conditional probability of informed trade is fairly compressed for short histories and expands monotonically with increasing trading history length. It is interesting that the median conditional probability falls with longer trading history, which is consistent with most trades not being informed. Panel (b) shows that an increasing fraction of trades are classified as likely-informed and more-likely-than-not informed as the trading history lengthens.

5.4. Two Case Studies

We now consider how the model classifies trades made by several prominent insiders accused of insider trading. We first consider Ken Lay and Jeff Skilling of Enron. Both executives sold large amounts of shares prior to Enron's bankruptcy filed in December 2001. Both men were found guilty of conspiracy and fraud. Skilling was found guilty of one count of insider trading and not guilty on nine other insider trading charges.

For each Enron executive, Figure 15 reports the trade-by-trade time-series of conditional probabilities (top figure of each panel) as well as the historical average abnormal return and associated standard error. Both Lay and Skilling started persistent selling in late 2000 that persisted through the first half of 2001. For the most part, the share price of Enron declined over this period, so these sales were quite profitable on average. As a result, the past average abnormal return increased dramatically for both men and its standard error tightened. The model thus classified trades made by the second quarter of 2001 as increasingly informed.

It is interesting to compare how the model classifies the trades of Lay and Skilling over this time period. Lay made over 200 trades and had a high average past abnormal return of about 3% by mid-2001. Skilling made fewer trades overall and started the episode with a lower past average abnormal return. Indeed, it was negative through most of the first quarter of 2001. The standard error of Skilling's past trades was also wider than that of Lay. As a result, the model classifies Skilling's later trades as more-likely-than-not informed, but does not classify them with the same confidence it does Lay's trades.

The second case study involves trading of Tesla by Elon Musk and Kimbal Musk in late 2021 and 2022. Elon Musk and Tesla directors (including Kimbal Musk) were sued in a shareholder derivative lawsuit by Tesla shareholder The Employees' Retirement System of Rhode Island (ERSRI). ERSRI alleged breach of fiduciary duties including unlawful stock sales in connection with Musk's purchase of Twitter. Figure 16 shows the evolution of the conditional probabilities for Elon and Kimbal Musk. There are stark differences in trading patterns and past profitability between the brothers. Post-2013, Elon Musk's trades have been quite profitable, generally exceeding 10% over that period. On the other hand, Kimbal Musk has consistently sold stock resulting in negative past average abnormal returns with relatively small standard errors. As a result, the model classifies Elon Musk's sales in late 2021 and 2022 as more-likely-than-not informed, while Kimbal Musk's sale in late 2021 has only about a 10% conditional probability of being informed.

6. Conclusion

Corporate insiders have access to private information about their firms. We use mixture model methods designed to account for noise in trading performance to assess the prevalence of informed trade by insiders required to disclose stock trades in US corporations. 30% of insiders make informed trades. Out-of-sample tests show that trades made by insiders most likely to have traded on information are predictive of future stock returns. Informed insider trade is most prevalent in chief financial officers and insider blockholders.

The fraction of informed insiders has remained fairly constant over time, but their information is impounded more quickly into prices in the most recent decade. This may be due to trading by sophisticated investors. We present evidence that these investors increased acquisition of insider trade disclosures from the SEC website over the same time period.

The insider-level model implies an insider-specific mixture model for trade returns that can be used to determine whether a trade was informed or not. The model results in an informed trade classification threshold that is customized based on an insider's return history. The model's conditional probability of an informed trade can be used by regulators to allocate enforcement resources devoted to monitoring corporate insider trade.

Appendix A. Derivation of $f_{\alpha^+|\bar{r}}$

The density of α_i conditional on an insider trading on information and having a realized average return $\bar{r}_i$ is denoted $f_{\alpha^+|\bar{r}}$. Here we show that $f_{\alpha^+|\bar{r}}$ is the density of a normal variable with mean of $\bar{r}_i - s_i^2/\mu$ and standard deviation of s_i that is truncated below at zero.

Using more general notation, we are interested in the distribution of x conditional on z, where $z = x + y$ with x distributed exponentially with mean $\mu = 1/\lambda$ and y distributed normally with mean zero and standard deviation s. First, note that using Bayes' rule, we can write

$$f_{x|z}(x|z) = \frac{f_x(x) f_{z|x}(z|x)}{f_z(z)} = \frac{f_x(x) f_y(z-x)}{f_z(z)} \,,$$

where the second equality follows from the fact that $f_{z|x}$ is just the distribution of y shifted by x. We plug in the functional forms of f_x, f_y, and the convolution f_z to obtain

$$f_{x|z}(x|z) = \frac{\mathbf{1}[x>0]\lambda \exp(-\lambda x) \frac{1}{s\sqrt{2\pi}} \exp\left(-\frac{(z-x)^2}{2s^2}\right)}{\int_{-\infty}^{\infty} \mathbf{1}[\tau>0]\lambda \exp(-\lambda\tau) \frac{1}{s\sqrt{2\pi}} \exp\left(-\frac{(z-\tau)^2}{2s^2}\right) d\tau} \,. \tag{A.1}$$

We want to show that this can be written in the form of the truncated normal density with normal mean $\tilde{m} = z - \lambda s^2$ and standard deviation s truncated below at 0; that is,

$$f_{x|z}(x|z) = \mathbf{1}[x>0] \cdot \frac{\phi(x;\tilde{m},s)}{1-\Phi(0;\tilde{m},s)} \,, \tag{A.2}$$

where $\phi(\cdot;m,s)$ and $\Phi(\cdot;m,s)$ are the density and cumulative distribution functions of a normal random variable with mean m and standard deviation s. With some algebra, the numerator of equation (A.1) can be written

$$\mathbf{1}[x>0]\lambda \exp(A) \frac{1}{s\sqrt{2\pi}} \exp\left(-\frac{(x-(z-\lambda s^2))^2}{2s^2}\right) = \mathbf{1}[x>0]\lambda \exp(A)\cdot\phi(x;\tilde{m},s) \tag{A.3}$$

where $A = \frac{1}{2}\lambda^2 s^2 - z\lambda$. Similarly, the denominator of equation (A.1) can be written

$$\lambda \exp(A) \int_0^{\infty} \frac{1}{s\sqrt{2\pi}} \exp\left(-\frac{(\tau-\tilde{m})^2}{2s^2}\right) d\tau = \lambda \exp(A) \left[1 - \Phi(0;\tilde{m},s)\right] \,. \tag{A.4}$$

Taking the ratio of equations (A.3) and (A.4) shows that (A.1) is equivalent to (A.2).

Appendix B. Normally-modified Truncated Normal Distribution

For insiders with an existing history of trades, the density of α_{ij} conditional on the insider trading on information is the convolution of a truncated normal distribution $f_{\alpha^+|\bar{r}_i}$ and a zero-mean normal variable:

$$\int_{-\infty}^{\infty} f_{\alpha^+|\bar{r}_i}(a) \cdot \phi(r_{ij} - a; \sigma_i)\ \mathrm{d}a\,.$$

Here we derive the analytical formula for this density.

Using more general notation, we are interested in the distribution of $z = y + x$ with x being a zero-mean normal random variable with standard deviation s and y being a normally distributed variable with mean μ and standard deviation σ that is truncated below at zero. We will show that the convolution of y and x is:

$$f_Z(z) = \int_{-\infty}^{\infty} f_Y(a) f_X(z-a)\ \mathrm{d}a = \frac{\exp\left(-\frac{(z-\mu)^2}{2(s^2+\sigma^2)}\right)}{\sqrt{2\pi(s^2+\sigma^2)}\Phi(\mu/\sigma)} \Phi\left(\frac{A}{\sigma\sqrt{B}}\right). \tag{B.1}$$

where $B = s^2/(s^2+\sigma^2)$ and $A = B\mu + (1-B)z$. Below denote ϕ and Φ as the standard normal density and cumulative distribution functions. Since y is truncated below at zero, we have

$$\begin{aligned} f_Z(z) &= \int_0^{\infty} f_Y(a) f_X(z-a)\ \mathrm{d}a \qquad\qquad (\mathrm{B.2}) \\ &= \int_0^{\infty} \frac{\phi(\frac{a-\mu}{\sigma})}{\sigma[1-\Phi(-\mu/\sigma)]} \cdot \frac{\phi(\frac{z-a}{s})}{s}\ \mathrm{d}a \\ &= \frac{1}{\sqrt{2\pi}\sigma s \Phi(\mu/\sigma)} \int_0^{\infty} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\left[\left(\frac{a-\mu}{\sigma}\right)^2 + \left(\frac{z-a}{s}\right)^2\right]\right)\ \mathrm{d}a\,. \end{aligned}$$

Expanding the term in brackets in the last line, we have

$$\begin{aligned} \left(\frac{a-\mu}{\sigma}\right)^2 + \left(\frac{z-a}{s}\right)^2 &= \frac{s^2}{s^2\sigma^2}\left(a^2 - 2a\mu + \mu^2\right) + \frac{\sigma^2}{s^2\sigma^2}\left(a^2 - 2az + z^2\right) \\ &= \frac{1}{s^2\sigma^2}\left[\left(s^2+\sigma^2\right)a^2 - 2s^2 a\mu - 2\sigma^2 az + s^2\mu^2 + \sigma^2 z^2\right] \\ &= \frac{s^2+\sigma^2}{s^2\sigma^2}\left[a^2 - 2aA + B\mu^2 + (1-B)z^2\right] \\ &= \frac{1}{\sigma^2 B}\left[(a-A)^2 - A^2 + B\mu^2 + (1-B)z^2\right]. \end{aligned}$$

Plugging in the expression for A, the last three terms can be reduced:

$$
\begin{aligned}
\frac{1}{\sigma^2 B}\Big[-A^2+B\mu^2+(1-B)z^2\Big] &= \frac{1}{\sigma^2 B}\Big[-\big(B^2\mu 2+(1-B)^2z^2+2B(1-B)\mu z\big)+B\mu^2+(1-b)z^2\Big] \\
&= \frac{1}{\sigma^2 B}\Big[\mu^2(B-B^2)-2B(1-B)\mu z+z^2\left(1-B-(1-B)^2\right)\Big] \\
&= \frac{1}{\sigma^2 B}\Big[B(1-B)(z-\mu)^2\Big] \\
&= \frac{(z-\mu)^2}{s^2+\sigma^2}\,.
\end{aligned}
$$

Thus, we can rewrite (B.2) to obtain the distribution in (B.1):

$$
\begin{aligned}
f_Z(z) &= \frac{\exp\left(-\frac{(z-\mu)^2}{2(s^2+\sigma^2)}\right)}{\sqrt{2\pi}\sigma s\Phi(\mu/\sigma)}\int_0^\infty \frac{1}{\sqrt{2\pi}}\exp\left(-\frac{1}{2}\left(\frac{a-A}{\sigma\sqrt{B}}\right)^2\right)\,\mathrm{d}a \\
&= \frac{\sqrt{B}\exp\left(-\frac{(z-\mu)^2}{2(s^2+\sigma^2)}\right)}{\sqrt{2\pi}s\Phi(\mu/\sigma)}\frac{1}{\sigma\sqrt{B}}\int_0^\infty \frac{1}{\sqrt{2\pi}}\exp\left(-\frac{1}{2}\left(\frac{a-A}{\sigma\sqrt{B}}\right)^2\right)\,\mathrm{d}a \\
&= \frac{\sqrt{B}\exp\left(-\frac{(z-\mu)^2}{2(s^2+\sigma^2)}\right)}{\sqrt{2\pi}s\Phi(\mu/\sigma)}\left(1-\Phi\left(\frac{-A}{\sigma\sqrt{B}}\right)\right) \\
&= \frac{\exp\left(-\frac{(z-\mu)^2}{2(s^2+\sigma^2)}\right)}{\sqrt{2\pi(s^2+\sigma^2)}\Phi(\mu/\sigma)}\Phi\left(\frac{A}{\sigma\sqrt{B}}\right)\,.
\end{aligned}
$$

References

Ahern, Kenneth R, 2017, Information networks: Evidence from illegal insider trading tips, *Journal of Financial Economics* 125, 26–47.

Akbas, Ferhat, Chao Jiang, and Paul D Koch, 2020, Insider investment horizon, *Journal of Finance* 75, 1579–1627.

Ali, Usman, and David Hirshleifer, 2017, Opportunism as a firm and managerial trait: Predicting insider trading profits and misconduct, *Journal of Financial Economics* 126, 490–515.

Biggerstaff, Lee, David Cicero, and M Babajide Wintoki, 2020, Insider trading patterns, *Journal of Corporate Finance* 64, 101654.

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

Chen, Huaizhi, Lauren Cohen, Umit Gurun, Dong Lou, and Christopher Malloy, 2020a, IQ from IP: Simplifying search in portfolio choice, *Journal of Financial Economics* 138, 118–137.

Chen, Yong, Michael Cliff, and Haibei Zhao, 2017, Hedge funds: The good, the bad, and the lucky, *Journal of Financial and Quantitative Analysis* 52, 1081–1109.

Chen, Yong, Bryan Kelly, and Wei Wu, 2020b, Sophisticated investors and market efficiency: Evidence from a natural experiment, *Journal of Financial Economics* 138, 316–341.

Cline, Brandon N, Sinan Gokkaya, and Xi Liu, 2017, The persistence of opportunistic insider trading, *Financial Management* 46, 919–964.

Cohen, Lauren, Christopher Malloy, and Lukasz Pomorski, 2012, Decoding inside information, *Journal of Finance* 67, 1009–1043.

Crane, Alan, and Kevin Crotty, 2020, How skilled are security analysts?, *Journal of Finance* 75, 1629–1675.

Crane, Alan, Kevin Crotty, and Tarik Umar, 2023, Hedge funds and public information acquisition, *Management Science* 69, 3241–3262.

Cziraki, Peter, and Jasmin Gider, 2021, The dollar profits to insider trading, *Review of Finance* 25, 1547–1580.

D'agostino, Ralph, and Egon S Pearson, 1973, Tests for departure from normality. empirical results for the distributions of b2 and $\sqrt{b1}$, *Biometrika* 60, 613–622.

DeMarzo, Peter M, Michael J Fishman, and Kathleen M Hagerty, 1998, The optimal enforcement of insider trading regulations, *Journal of Political Economy* 106, 602–632.

Dim, Chukwuma, 2023, Social media analysts' skills: Insights from text-implied beliefs, Working paper.

Dye, Ronald A, 1984, Inside trading and incentives, *Journal of Business* 295–313.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

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

Goldie, Brad, Chao Jiang, Paul Koch, and M Babajide Wintoki, 2023, Indirect insider trading, *Journal of Financial and Quantitative Analysis* 58, 2327–2364.

Harvey, Campbell R, and Yan Liu, 2018, Detecting repeatable performance, *The Review of Financial Studies* 31, 2499–2552.

Hirshleifer, Jack, 1971, The private and social value of information and the reward to inventive activity, *American Economic Review* 61, 561–574.

Huddart, Steven, John S Hughes, and Carolyn B Levine, 2001, Public disclosure and dissimulation of insider trades, *Econometrica* 69, 665–681.

Jaffe, Jeffrey F, 1974, Special information and insider trading, *The Journal of Business* 47, 410–428.

Jeng, Leslie A, Andrew Metrick, and Richard Zeckhauser, 2003, Estimating the returns to insider trading: A performance-evaluation perspective, *Review of Economics and Statistics* 85, 453–471.

Kacperczyk, Marcin, and Emiliano S Pagnotta, 2019, Chasing private information, *The Review of Financial Studies* 32, 4997–5047.

Kakhbod, Ali, Seyed Mohammad Kazempour, Dmitry Livdan, and Norman Schürhoff, 2023, Finfluencers, Working paper.

Koppell, Jonathan, 2011, *Origins of shareholder advocacy* (Springer).

Leland, Hayne E, 1992, Insider trading: Should it be prohibited?, *Journal of Political Economy* 100, 859–887.

McLean, R David, and Jeffrey Pontiff, 2016, Does academic research destroy stock return predictability?, *The Journal of Finance* 71, 5–32.

Meulbroek, Lisa K, 1992, An empirical analysis of illegal insider trading, *The Journal of Finance* 47, 1661–1699.

Ravina, Enrichetta, and Paola Sapienza, 2010, What do independent directors know? evidence from their trading, *The Review of Financial Studies* 23, 962–1003.

Seyhun, H Nejat, 1986, Insiders' profits, costs of trading, and market efficiency, *Journal of Financial Economics* 16, 189–212.

Seyhun, H Nejat, 1992a, The effectiveness of the insider-trading sanctions, *The Journal of Law and Economics* 35, 149–182.

Seyhun, H Nejat, 1992b, Why does aggregate insider trading predict future stock returns?, *The Quarterly Journal of Economics* 107, 1303–1331.

Wang, Weimin, Yong-Chul Shin, and Bill B Francis, 2012, Are CFOs' trades more informative than ceos' trades?, *Journal of Financial and Quantitative Analysis* 47, 743–762.

Table 1: Summary Statistics of Insider Trading History

This table reports cross-sectional distributional statistics of average 21-trading day abnormal returns and standard errors following trades made by corporate insiders. Average abnormal returns are defined in Equation (9) and are the average of long positions in purchased stocks and short positions in sold stocks, relative to the market benchmark. The sample contains all insiders who traded on at least 10 distinct days from 1985 to 2022. Because mixture model methods can be sensitive to outliers, the sample is trimmed at the 1% and 99% level of average abnormal returns. The table also reports the distribution of the number of distinct trading days. The table reports the fraction of insiders with sample average abnormal returns that are (1) positive, (2) significantly positive at the 5% level, (3) significantly positive at a 10% level (both in a one-sided test). The table also reports the p-value from a test of normality of the average abnormal return distribution following D'agostino and Pearson (1973).

	Average Abnormal Return	Standard Error	#(Trades)
Mean	0.0066	0.0248	31
SD	0.0468	0.0175	49
P1	-0.1226	0.0050	10
P10	-0.0428	0.0094	11
P25	-0.0162	0.0133	13
P50	0.0037	0.0200	19
P75	0.0268	0.0309	31
P90	0.0604	0.0457	57
P99	0.1550	0.0898	201
Skewness	0.40	2.49	15
Excess Kurtosis	2.47	11.80	504
Fraction positive	0.55		
Significant 5%	0.18		
Significant 10%	0.24		
Normality p-value	0.0000		
N	54274		

Table 2: A Mixture Model of the Cross-Section of Insiders

This table reports mixture model parameter estimates for the cross-section of corporate insider average abnormal returns. Equation (3) is numerically maximized in π and μ. To limit the effect of outliers, the sample is first trimmed at the 1 and 99% percentiles of average abnormal returns. The point estimates, negative log-likelihood, and number of observations in the trimmed sample are reported. The reported confidence interval for each parameter is bootstrapped. Specifically, the model is estimated on 1,000 bootstrapped samples (each is formed by sampling with replacement). The reported confidence intervals are the 1st and 99th percentiles of the bootstrapped parameter estimates.

	π		μ
Parameter Estimate	0.2855		0.0360
Confidence Interval:			
Lower	0.2775		0.0351
Upper	0.2943		0.0370
Negative Log-Likelihood		-84,209.61	
Observations		54,274	

Table 3: The Insider's Role in the Firm Relates to their Information.

This table reports regressions of conditional probabilities (columns 1 and 2) and expectations (Columns 3 and 4) of informed insider trading on their role(s) within the firm as disclosed on the Form 4 filing. The mixture model is estimated in an expanding window fashion each year using the latest average abnormal returns and standard errors for each insider with at least ten trades prior to that year's end. Based on the estimated parameters and each insider's average abnormal return and standard error, a conditional probability that the insider engages in informed trade and the conditional expectation of insider informed trade is calculated. Expectations are reported in percent. We include a dummy variable for each of the roles disclosed by the insider. Columns 1 and 3 include indicators for any role in the aggregate category displayed. Columns 2 and 4 estimate the model on indicators for all roles but only display coefficients for roles representing at least 5% of the sample. Standard errors are clustered by insider and year. t-statistics are reported below coefficient estimates, and statistical significance is represented by * ($p < 0.10$), ** ($p < 0.05$), and *** ($p < 0.01$).

	(1) $\tilde{\pi}$	(2) $\tilde{\pi}$	(3) $\tilde{\alpha}$	(4) $\tilde{\alpha}$
C-Suite Executives	0.03***		0.23***	
	(7.23)		(6.37)	
Directors	0.00		-0.02	
	(1.17)		(-0.74)	
Insider Owners	0.07***		0.83***	
	(12.41)		(15.48)	
Non-C-Suite Officers	-0.02***		-0.17***	
	(-5.98)		(-4.84)	
Non-officer managers	0.00		0.04	
	(0.25)		(0.44)	
CEO		0.00		-0.03
		(0.58)		(-0.53)
CFO		0.03***		0.26***
		(6.17)		(4.74)
Inside Block > 10%		0.07***		0.93***
		(10.53)		(12.46)
Chairman		-0.01		-0.11**
		(-1.53)		(-2.09)
Director		-0.01***		-0.15***
		(-3.49)		(-3.38)
President		0.01*		0.12**
		(1.72)		(2.31)
Executive VP		-0.03***		-0.18***
		(-5.21)		(-3.61)
Senior VP		-0.03***		-0.25***
		(-5.62)		(-4.96)
VP		-0.00		0.05
		(-0.58)		(0.90)
Officer		-0.03***		-0.21***
		(-5.34)		(-4.44)
Divisional Officer		-0.01***		-0.14***
		(-3.03)		(-3.13)
Officer and Director		0.00		0.08
		(0.44)		(1.39)
Observations	966,445	966,445	966,445	966,445
Adjusted R^2	0.012	0.018	0.017	0.024

Table 4: Trades by Informed Insiders Predict Future Returns–Regression

This table reports regressions of monthly stock returns as a function of insider buying and selling activity in the prior month. The mixture model is estimated in an expanding window fashion each year using the latest average abnormal returns and standard errors for each insider with at least ten trades prior to that year's end. Based on the estimated parameters and each insider's average abnormal return and standard error, the conditional expectation of insider informed trade is calculated. Insiders are sorted into quintiles by the conditional expectation. Buy Quintile 5 (Sell Quintile 5) is an indicator variable for whether any insider in the top quintile bought (sold) shares in month t. The other quintile indicators are similarly defined. Control variables include size, book-to-market, returns in month $t-1$, and months $t-12$ to $t-2$. Month-fixed effects are included in even-numbered columns. Standard errors are clustered by firm and month. t-statistics are reported below coefficient estimates, and statistical significance is represented by * ($p < 0.10$), ** ($p < 0.05$), and *** ($p < 0.01$). p values for tests of whether the coefficients on the High and Low quintiles differ are reported in the table footer.

	Dependent Variable: Return$_{t+1}$					
	(1)	(2)	(3)	(4)	(5)	(6)
Buy Quintile 1 (Lo)	0.12	0.18			0.07	0.11
	(0.62)	(0.86)			(0.34)	(0.53)
Buy Quintile 2	0.49**	0.40**			0.45**	0.34**
	(2.32)	(2.33)			(2.08)	(2.02)
Buy Quintile 3	0.57**	0.43***			0.52*	0.37**
	(2.19)	(2.89)			(1.92)	(2.46)
Buy Quintile 4	0.87***	0.76***			0.83**	0.69***
	(2.92)	(4.02)			(2.58)	(3.46)
Buy Quintile 5 (Hi)	0.93***	0.86***			0.88***	0.79***
	(3.53)	(5.19)			(3.18)	(4.90)
Sell Quintile 1 (Lo)			-0.07	-0.14	0.17	0.07
			(-0.48)	(-1.50)	(1.28)	(0.66)
Sell Quintile 2			-0.19	-0.17**	0.02	0.00
			(-1.55)	(-2.01)	(0.17)	(0.01)
Sell Quintile 3			-0.36***	-0.32***	-0.15	-0.15
			(-2.82)	(-3.04)	(-1.18)	(-1.48)
Sell Quintile 4			-0.36***	-0.38***	-0.14	-0.19*
			(-2.84)	(-3.39)	(-1.18)	(-1.92)
Sell Quintile 5 (Hi)			-0.52***	-0.44***	-0.27*	-0.23**
			(-3.67)	(-3.64)	(-1.86)	(-2.00)
Size	-0.12	-0.07	-0.15*	-0.09	-0.13*	-0.07
	(-1.50)	(-1.14)	(-1.95)	(-1.62)	(-1.66)	(-1.27)
BM	0.41**	0.30**	0.43**	0.30**	0.41**	0.29**
	(2.17)	(2.15)	(2.17)	(2.14)	(2.14)	(2.12)
Ret(t-1)	-0.39	0.27	-0.50	0.21	-0.34	0.32
	(-0.19)	(0.24)	(-0.24)	(0.18)	(-0.16)	(0.29)
Ret(t-12,t-2)	0.35	0.60**	0.33	0.59**	0.36	0.61**
	(0.77)	(2.07)	(0.74)	(2.04)	(0.79)	(2.09)
Constant	3.73**		4.87***		4.02**	
	(2.31)		(3.16)		(2.52)	
Time FE	N	Y	N	Y	N	Y
Adj R^2	0.0028	0.1398	0.0025	0.1395	0.0029	0.1398
Observations	180715	180715	180715	180715	180715	180715
p(Buy Hi-Lo)	0.0016	0.0060			0.0015	0.0059
p(Sell Hi-Lo)			0.0104	0.0355	0.0115	0.0381

Table 5: Trades by Informed Insiders Predict Future Returns–Portfolios

This table reports returns for portfolios formed by sorting on (1) an insider's conditional expectation $\tilde{\alpha}$ and (2) their signed trading activity. The mixture model is estimated in an expanding window fashion each year using the latest average abnormal returns and standard errors for each insider with at least ten trades prior to that year's end. Based on the estimated parameters and each insider's average abnormal return and standard error, the conditional expectation of insider informed trade is calculated. Insiders are sorted into quintiles by the conditional expectation. This sorting is done based on an insider's trade history up to the prior year's end for a given month's portfolio formation. The second sort is based on signed insider order flow. Specifically, for each insider-stock pair, the signed order flow is calculated for a given month as a percentage of the stock's total trading volume that month. In a given month's portfolio formation, quintiles of signed order flow are formed by sorting insider-stock order flows. Portfolios are equal-weighted by insider-stock observations within each of the 25 quintile combinations. Panel A reports the average excess returns of each portfolio. Panels B, C, and D report alphas using market, Fama-French-Carhart, and Fama-French Five Factor + momentum benchmarks, respectively.

Panel A: Average Excess Returns

	Conditional Expectation					
	Low	2	3	4	High	Hi-Lo
Strong Sells	0.0110	0.0112	0.0044	0.0076	0.0074	-0.0036
	(4.12)	(3.84)	(1.41)	(2.33)	(2.07)	(-1.70)
2	0.0104	0.0087	0.0079	0.0066	0.0092	-0.0012
	(3.44)	(2.76)	(2.31)	(1.88)	(2.63)	(-0.61)
3	0.0071	0.0081	0.0060	0.0047	0.0087	0.0014
	(2.35)	(2.58)	(1.76)	(1.36)	(2.27)	(0.61)
4	0.0125	0.0115	0.0143	0.0133	0.0125	0.0000
	(4.54)	(3.96)	(4.38)	(3.84)	(3.47)	(0.00)
Strong Buys	0.0117	0.0175	0.0180	0.0202	0.0213	0.0096
	(4.55)	(5.90)	(5.47)	(6.39)	(6.42)	(4.49)
SB Minus SS	0.0007	0.0063	0.0136	0.0127	0.0139	0.0132
	(0.35)	(2.53)	(5.05)	(4.96)	(5.13)	(4.65)

Panel B: CAPM Alpha

	Conditional Expectation					
	Low	2	3	4	High	Hi-Lo
Strong Sells	0.0103	0.0102	0.0034	0.0064	0.0057	-0.0046
	(3.42)	(3.16)	(0.99)	(1.80)	(1.47)	(-2.12)
2	0.0100	0.0083	0.0077	0.0058	0.0082	-0.0018
	(2.98)	(2.37)	(2.04)	(1.53)	(2.15)	(-0.90)
3	0.0063	0.0075	0.0057	0.0043	0.0080	0.0016
	(1.90)	(2.15)	(1.50)	(1.11)	(1.93)	(0.65)
4	0.0118	0.0108	0.0134	0.0122	0.0120	0.0003
	(3.91)	(3.47)	(3.81)	(3.24)	(3.19)	(0.11)
Strong Buys	0.0100	0.0161	0.0162	0.0182	0.0195	0.0094
	(3.66)	(5.07)	(4.59)	(5.40)	(5.49)	(4.41)
SB Minus SS	-0.0002	0.0058	0.0127	0.0118	0.0138	0.0140
	(-0.11)	(2.27)	(4.52)	(4.46)	(5.05)	(4.96)

Table 5 (continued)

Panel C: Fama-French Three Factor + Momentum Alpha

	Conditional Expectation					
	Low	2	3	4	High	Hi-Lo
Strong Sells	0.0111	0.0108	0.0045	0.0074	0.0069	-0.0042
	(3.44)	(3.05)	(1.19)	(1.97)	(1.65)	(-1.85)
2	0.0111	0.0092	0.0087	0.0078	0.0094	-0.0017
	(3.06)	(2.38)	(2.10)	(1.87)	(2.15)	(-0.80)
3	0.0072	0.0086	0.0068	0.0057	0.0085	0.0011
	(1.90)	(2.14)	(1.58)	(1.32)	(1.89)	(0.45)
4	0.0122	0.0114	0.0136	0.0131	0.0136	0.0013
	(3.78)	(3.44)	(3.59)	(3.28)	(3.38)	(0.54)
Strong Buys	0.0104	0.0160	0.0167	0.0185	0.0202	0.0099
	(3.52)	(4.70)	(4.37)	(4.97)	(5.27)	(4.54)
SB Minus SS	-0.0007	0.0052	0.0122	0.0111	0.0133	0.0140
	(-0.35)	(1.98)	(4.19)	(4.08)	(4.73)	(5.02)

Panel D: Fama-French Five Factor + Momentum Alpha

	Conditional Expectation					
	Low	2	3	4	High	Hi-Lo
Strong Sells	0.0116	0.0121	0.0054	0.0082	0.0081	-0.0035
	(3.51)	(3.32)	(1.38)	(2.09)	(1.87)	(-1.50)
2	0.0121	0.0102	0.0099	0.0094	0.0110	-0.0011
	(3.27)	(2.57)	(2.29)	(2.17)	(2.40)	(-0.47)
3	0.0083	0.0099	0.0084	0.0069	0.0099	0.0014
	(2.10)	(2.38)	(1.85)	(1.53)	(2.14)	(0.59)
4	0.0131	0.0116	0.0137	0.0142	0.0150	0.0020
	(3.87)	(3.29)	(3.48)	(3.43)	(3.58)	(0.80)
Strong Buys	0.0108	0.0169	0.0182	0.0188	0.0214	0.0106
	(3.42)	(4.65)	(4.40)	(4.74)	(5.16)	(4.60)
SB Minus SS	-0.0008	0.0048	0.0128	0.0106	0.0133	0.0141
	(-0.39)	(1.76)	(4.41)	(3.73)	(4.51)	(4.79)

Table 6: Post-2012 Daily Portfolio Analysis of Informed Insiders

This table reports post-2012 returns for portfolios formed by sorting on (1) an insider's conditional expectation $\tilde{\pi}$ and (2) whether the trade is a buy or a sell. The mixture model is estimated in an expanding window fashion each year using the latest average abnormal returns and standard errors for each insider with at least ten trades prior to that year's end. Based on the estimated parameters and each insider's average abnormal return and standard error, the conditional expectation of insider informed trade is calculated. Insiders are sorted into quintiles by the conditional expectation. This sorting is done based on an insider's trade history up to the prior year's end for a given date's portfolio formation. The second sort is based on whether the trade is a buy or a sell. A stock from a given trade enters the buy or sell portfolio the day following the trade and is included for forty days. The portfolio holds all stocks included in the portfolio that day at equal weights. The portfolio returns are converted to monthly returns for comparability with the monthly portfolio analysis. Panel A reports the average excess returns of each portfolio. Panels B, C, and D report alphas using market, Fama-French-Carhart, and Fama-French Five Factor + momentum benchmarks, respectively.

Panel A: Average Excess Returns

	Conditional Expectation					
	Low	2	3	4	High	Hi-Lo
Sells	0.0112	0.0105	0.0094	0.0097	0.0090	-0.0022
	(2.31)	(2.07)	(1.75)	(1.83)	(1.66)	(-1.32)
Buys	0.0176	0.0198	0.0211	0.0211	0.0243	0.0067
	(3.69)	(3.82)	(3.98)	(4.02)	(4.38)	(3.05)
Buys Minus Sells	0.0065	0.0093	0.0117	0.0114	0.0153	0.0088
	(2.71)	(3.41)	(4.30)	(4.49)	(6.40)	(3.98)

Panel B: CAPM Alpha

	Conditional Expectation					
	Low	2	3	4	High	Hi-Lo
Sells	-0.0006	-0.0017	-0.0033	-0.0028	-0.0035	-0.0028
	(-0.46)	(-1.05)	(-1.68)	(-1.35)	(-1.49)	(-1.77)
Buys	0.0074	0.0090	0.0100	0.0101	0.0127	0.0053
	(2.81)	(2.99)	(3.29)	(3.35)	(4.00)	(2.50)
Buys Minus Sells	0.0080	0.0107	0.0133	0.0129	0.0161	0.0081
	(3.39)	(3.94)	(4.95)	(5.10)	(6.68)	(3.69)

Table 6 (continued)

Panel C: Fama-French Three Factor + Momentum Alpha

	Conditional Expectation					
	Low	2	3	4	High	Hi-Lo
Sells	0.0001	-0.0007	-0.0019	-0.0014	-0.0019	-0.0021
	(0.17)	(-0.77)	(-1.86)	(-1.23)	(-1.42)	(-1.51)
Buys	0.0087	0.0105	0.0117	0.0118	0.0146	0.0059
	(5.02)	(5.57)	(5.84)	(5.71)	(6.59)	(3.03)
Buys Minus Sells	0.0086	0.0112	0.0137	0.0132	0.0166	0.0080
	(5.00)	(5.86)	(6.84)	(6.44)	(7.61)	(3.74)

Panel D: Fama-French Five Factor + Momentum Alpha

	Conditional Expectation					
	Low	2	3	4	High	Hi-Lo
Sells	0.0007	0.0001	-0.0009	-0.0005	-0.0010	-0.0017
	(0.87)	(0.10)	(-1.01)	(-0.53)	(-0.81)	(-1.29)
Buys	0.0092	0.0111	0.0124	0.0128	0.0158	0.0066
	(5.35)	(5.90)	(6.26)	(6.33)	(7.45)	(3.51)
Buys Minus Sells	0.0085	0.0110	0.0134	0.0133	0.0168	0.0083
	(4.95)	(5.75)	(6.72)	(6.48)	(7.74)	(3.89)

Table 7: Conditional Informed Insider Measures and Existing Measures

This table reports regressions of conditional probabilities (Panel A) and expectations (Panel B) of informed insider trading on other classifiers of informed insider trading. The mixture model is estimated in an expanding window fashion each year using the latest average abnormal returns and standard errors for each insider with at least ten trades prior to that year's end. Based on the estimated parameters and each insider's average abnormal return and standard error, a conditional probability that the insider engages in informed trade and the conditional expectation of insider informed trade is calculated. Routine insiders are calculated following Cohen, Malloy, and Pomorski (2012). Investor horizon is calculated following Akbas, Jiang, and Koch (2020). High QEA Profitability represents the top quintile of insider profits ahead of quarterly earnings announcements and is calculated following Ali and Hirshleifer (2017). Standard errors are clustered by insider and year. t-statistics are reported below coefficient estimates, and statistical significance is represented by * ($p < 0.10$), ** ($p < 0.05$), and *** ($p < 0.01$).

Panel A

	Dependent Variable: Conditional Probability ($\tilde{\pi}$)		
	(1)	(2)	(3)
Non-Routine	0.04^{***}		
	(15.07)		
Medium Horizon		0.04^{***}	
		(12.10)	
Short Horizon		0.09^{***}	
		(15.82)	
High QEA Profitability			0.05^{***}
			(9.13)
Constant	0.22^{***}	0.22^{***}	0.25^{***}
	(47.88)	(64.59)	(40.10)
Adj R^2	0.0066	0.0211	0.0053
Observations	130671	111582	24639

Panel B

	Dependent Variable: Conditional Expectation ($\tilde{\alpha}$)		
	(1)	(2)	(3)
Non-Routine	0.0029^{***}		
	(17.05)		
Medium Horizon		0.0024^{***}	
		(12.17)	
Short Horizon		0.0057^{***}	
		(20.60)	
High QEA Profitability			0.0033^{***}
			(9.07)
Constant	0.0060^{***}	0.0057^{***}	0.0077^{***}
	(20.35)	(29.74)	(20.61)
Adj R^2	0.0061	0.0227	0.0058
Observations	130671	111582	24639

Table 8: Trades by Informed Insiders Predict Future Returns Controlling for Existing Measures

This table reports regressions of monthly stock returns as a function of insider buying and selling activity in the prior month. The mixture model is estimated in an expanding window fashion each year using the latest average abnormal returns and standard errors for each insider with at least ten trades prior to that year's end. Based on the estimated parameters and each insider's average abnormal return and standard error, the conditional expectation of insider informed trade is calculated. Insiders are sorted into quintiles on the basis of the conditional expectation. Buy Quintile 5 (Sell Quintile 5) is an indicator variable for whether any insider in the top quintile bought (sold) shares in month t. The other quintile indicators are similarly defined. Similar indicator variables are calculated for buys and sells made by three sets of insiders: (1) routine and non-routine insiders (Cohen et al., 2012); (2) long, medium, and short horizon insiders (Akbas et al., 2020); and (3) the highest quintile of QEA Profitability (Ali and Hirshleifer, 2017). Control variables include size, book-to-market, returns in month $t-1$, and months $t-12$ to $t-2$. Month-fixed effects and controls are included in all columns. Standard errors are clustered by firm and month. t-statistics are reported below coefficient estimates, and statistical significance is represented by * ($p < 0.10$), ** ($p < 0.05$), and *** ($p < 0.01$). p values for tests of whether the coefficients on the High and Low quintiles differ are reported in the table footer.

	Dependent Variable: Return_{t+1}					
	(1)	(2)	(3)	(4)	(5)	(6)
Buy Quintile 1 (Lo)		0.23		0.17		0.10
		(1.17)		(0.83)		(0.51)
Buy Quintile 2		0.43**		0.37**		0.33**
		(2.54)		(2.29)		(2.00)
Buy Quintile 3		0.44***		0.40**		0.37**
		(2.67)		(2.55)		(2.44)
Buy Quintile 4		0.75***		0.72***		0.69***
		(3.35)		(3.48)		(3.44)
Buy Quintile 5 (Hi)		0.85***		0.79***		0.78***
		(4.94)		(4.50)		(4.78)
Sell Quintile 1 (Lo)		0.02		0.05		0.06
		(0.24)		(0.51)		(0.61)
Sell Quintile 2		-0.02		0.00		-0.00
		(-0.27)		(0.00)		(-0.04)
Sell Quintile 3		-0.16*		-0.14		-0.15
		(-1.65)		(-1.27)		(-1.53)
Sell Quintile 4		-0.21**		-0.18		-0.20*
		(-2.08)		(-1.61)		(-1.96)
Sell Quintile 5 (Hi)		-0.24**		-0.21		-0.23**
		(-2.12)		(-1.64)		(-2.03)
Nonroutine Buy	0.45***	-0.07				
	(2.92)	(-0.45)				
Nonroutine Sell	-0.17**	0.02				
	(-2.30)	(0.20)				
Routine Buy	-0.26	-0.70***				
	(-1.06)	(-2.74)				
Routine Sell	0.19	0.29**				
	(1.52)	(2.13)				
Long Horizon Buy			0.04	-0.36*		
			(0.20)	(-1.88)		
Med Horizon Buy			0.49**	0.10		
			(2.14)	(0.48)		
Short Horizon Buy			0.86***	0.45**		
			(3.63)	(2.00)		
Long Horizon Sell			-0.00	0.15		
			(-0.05)	(1.42)		
Medium Horizon Sell			-0.36***	-0.22**		
			(-4.10)	(-2.10)		
Short Horizon Sell			-0.37***	-0.18		
			(-3.43)	(-1.50)		
QEA Profitability Q5 Buy					0.64	0.19
					(1.60)	(0.46)
QEA Profitability Q5 Sell					-0.03	0.14
					(-0.15)	(0.69)
Time FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Adj R^2	0.1394	0.1399	0.1395	0.1399	0.1393	0.1398
Observations	180715	180715	180715	180715	180715	180715
p(Buy Hi-Lo)		0.0116		0.0104		0.0064
p(Sell Hi-Lo)		0.0570		0.0593		0.0383

Table 9: Incorporating Existing Proxies

This table reports mixture model parameter estimates for the cross-section of corporate insider average abnormal returns. The model parameterizes π as a function of whether the indicated empirical proxy either classifies an insider as one who uses informed, one who does not use information, or does not classify the insider due to the insider not satisfying sample screens (Equation 10). To limit the effect of outliers, the sample is first trimmed at the 1 and 99% percentiles of average abnormal returns. The point estimates, negative log-likelihood, and number of observations in the trimmed sample are reported. The reported confidence interval for each parameter is bootstrapped. Specifically, the model is estimated on 1,000 bootstrapped samples (each is formed by sampling with replacement). The reported confidence intervals are the 1st and 99th percentiles of the bootstrapped parameter estimates.

Panel A. Informed: Non-Routine (Cohen, Malloy, and Pomorski, 2012)

	π			
	Unclassified	Uninformed	Informed	μ
Parameter Estimate	0.4395	0.1365	0.2368	0.0368
Confidence Interval				
Lower	0.4245	0.1141	0.2270	0.0358
Upper	0.4557	0.1594	0.2470	0.0377
Fraction of Data	0.2539	0.0708	0.6752	
Negative Log-Likelihood		-84,631.38		
Observations		54,274		

Panel B. Informed: High QEA Profitability (Ali and Hirshleifer, 2017)

	π			
	Unclassified	Uninformed	Informed	μ
Parameter Estimate	0.2970	0.1992	0.3329	0.0360
Confidence Interval				
Lower	0.2879	0.1829	0.3052	0.0351
Upper	0.3071	0.2179	0.3607	0.0369
Fraction of Data	0.7974	0.1294	0.0731	
Negative Log-Likelihood		-84,274.75		
Observations		54,274		

Panel C. Informed: Short Horizon (Akbas, Jiang, and Koch, 2020)

	π			
	Unclassified	Uninformed	Informed	μ
Parameter Estimate	0.3661	0.1447	0.4029	0.0366
Confidence Interval				
Lower	0.3553	0.1342	0.3699	0.0356
Upper	0.3775	0.1559	0.4337	0.0375
Fraction of Data	0.5595	0.3710	0.0794	
Negative Log-Likelihood		-84,788.79		
Observations		54,274		

Table 10: Conditional Probability that an Individual Trade is Informed

This table reports summary statistics of the conditional probability that a trade is informed (15) for individual insider trades. The first column reports the sample average of the conditional probability τ. Columns 2 through 4 report the sample average of indicators for whether the conditional probability is less than 5% (Likely Uninformed), greater than 50% (More-likely-than-not Informed), and greater than 95% (Likely Informed), respectively. The statistics are reported separately for the overall sample, for purchases and sales, for quintiles of trading strength, and conditional on the number of past trades made by the insider. Trading strength is signed trading volume divided by the stock's monthly volume. Trading strength quintiles are formed monthly.

	Average Conditional Probability τ	Fraction Likely Uninformed $1(\tau < 0.05)$	Fraction More-likely-than-not Informed $1(\tau > 0.5)$	Fraction Likely Informed $1(\tau > 0.95)$
Overall	0.2978	0.0883	0.1657	0.0672
Trade Direction				
Purchases	0.3499	0.0722	0.2196	0.1017
Sales	0.2735	0.0959	0.1401	0.0509
Trade Strength				
Strong Sell	0.3066	0.0902	0.1813	0.0719
2	0.2731	0.0998	0.1406	0.0495
3	0.2651	0.0896	0.1253	0.0452
4	0.2858	0.0830	0.1444	0.0566
Strong Buy	0.3590	0.0799	0.2388	0.1143
#(Past Trades)				
Less than 10	0.2990	0.0073	0.1148	0.0301
10-25	0.2842	0.0724	0.1697	0.0598
26-50	0.2858	0.1232	0.1968	0.0835
More than 50	0.3208	0.2947	0.2708	0.1617

Figure 1: Distributions of True and Estimated Informed Insider Trading

This figure illustrates the mixture method of informed insider trading. Panel (a) shows the relative frequencies of true informed insider trading. A fraction π of insiders trade on information that is exponentially distributed with mean μ (the hatched purple bins). The remaining $1 - \pi$ insiders do not trade on information (grey bins). Panel (b) shows the relative frequencies of estimated abnormal returns for insiders that exploit private information (hatched purple), insiders that do not (grey bins), and the unconditional distribution (black line). Estimated abnormal returns exhibit additional variation due to error in estimating true informed trading, resulting in more dispersed distributions in Panel (b) than in Panel (a). The parameter values for this example are $\pi = 0.7$, $\mu = 0.025$, and a standard error $s_i = 0.015$ for all insiders.

(a) True Informed Insider Trading

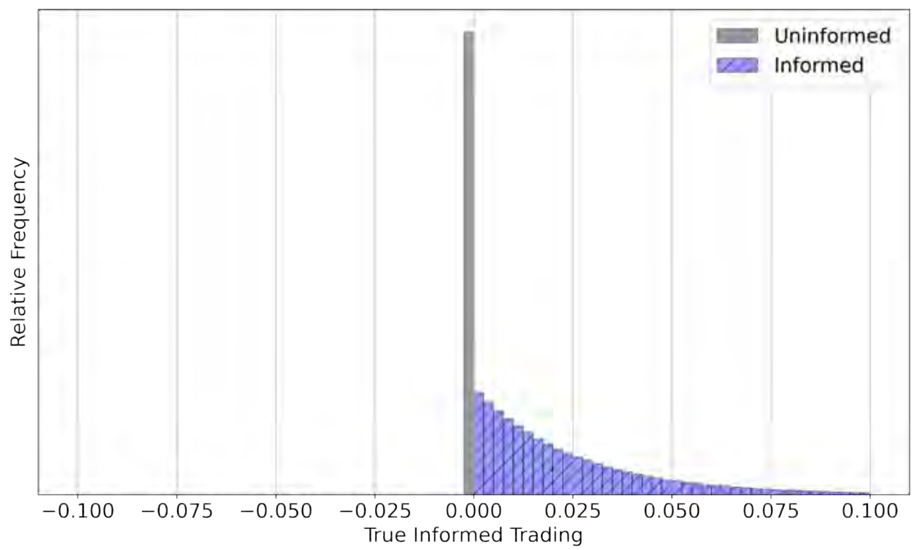

(b) Estimated Abnormal Return

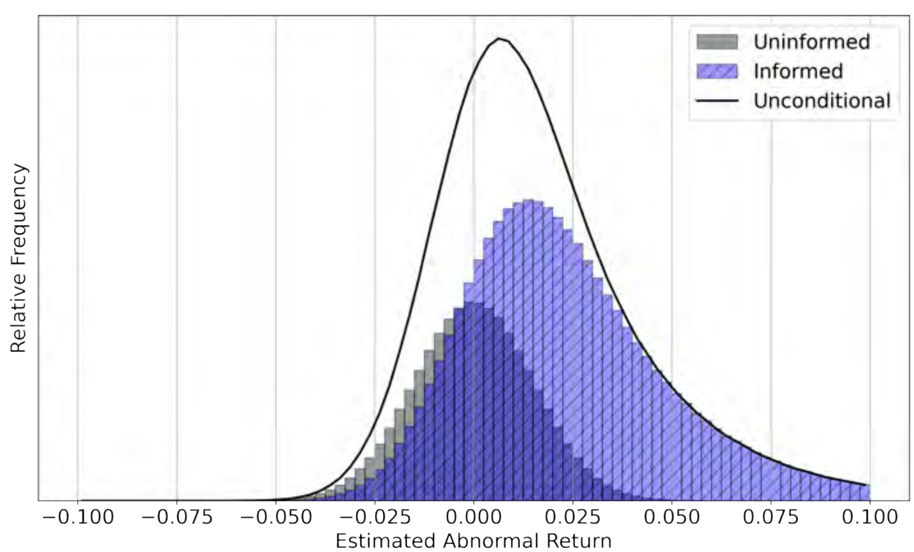

Figure 2: Conditional Probabilities and Expectations
This figure illustrates conditional probabilities and expectations in the mixture model method of informed insider trading as a function of the estimated average abnormal return and its standard error (i.e. its noise). Panel (a) shows the probability an insider trades on information conditional on their average abnormal return and its standard error. Panel (b) shows the conditional expectation of an insider's information, conditional on their average abnormal return and its standard error. The parameter values for this example are $\pi = 0.7$, $\mu = 0.025$, and the standard errors (noise) indicated in the legend.

(a) Conditional Probability Insider is Informed

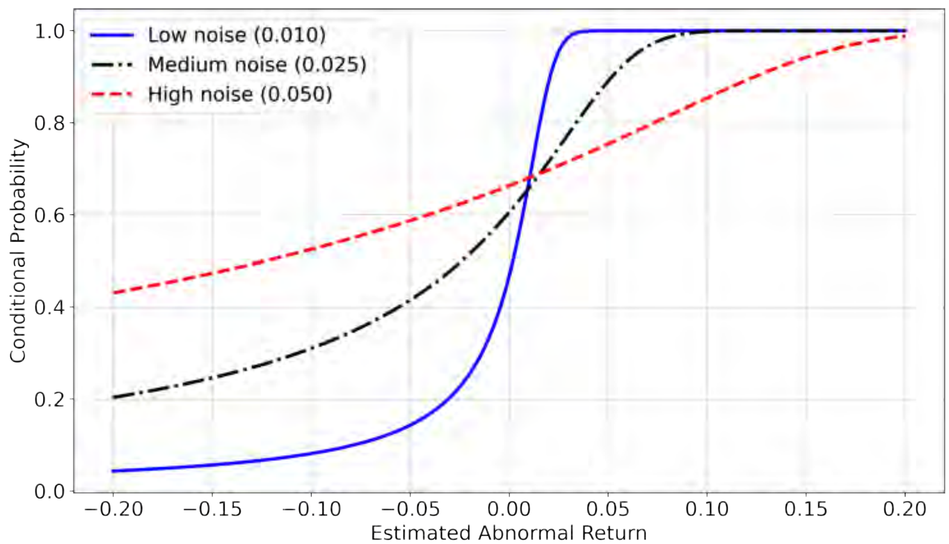

(b) Conditional Expectation of Insider Information

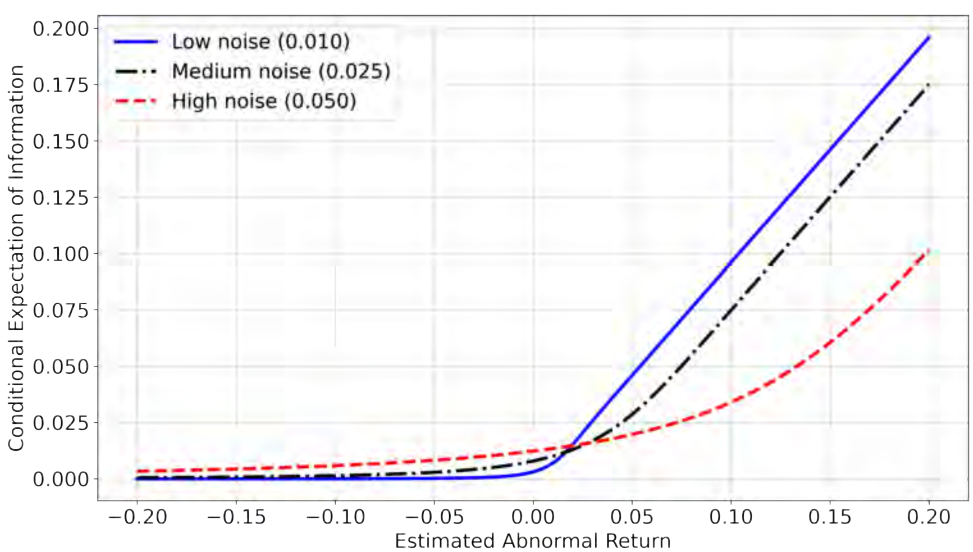

Figure 6: Cumulative Returns (Monthly)

This figure plots cumulative returns for portfolios formed by sorting on (1) an insider's conditional expectation $\tilde{\alpha}$ and (2) their signed trading activity, as in Table 5. The mixture model is estimated in an expanding window fashion each year using the latest average abnormal return and standard error for each insider with at least ten trades prior to that year's end. Based on the estimated parameters and each insider's average abnormal return and standard error, the conditional expectation of insider informed trade is calculated. Insiders are sorted into quintiles on the basis of the conditional expectation. This sorting is done based on an insider's trade history up to the prior year end for a given month's portfolio formation. The second sort is based on signed insider order flow. Specifically, for each insider-stock pair, the signed order flow is calculated for a given month as a percentage of the stock's total trading volume that month. In a given month's portfolio formation, quintiles of signed order flow are formed by sorting insider-stock order flows. Portfolios are equal-weighted by insider-stock observations within each of the 25 quintile combinations. Panel A reports cumulative returns for hedge portfolios that buy stocks with strong inside buying pressure and sell stocks with strong inside selling pressure. The black solid (red dashed) line represents this strategy for insiders in the top (bottom) quintile of ex-ante conditional expectation. Panel B reports the cumulative performance for hedge portfolios that either (1) buy the top $\tilde{\alpha}$ quintile's strong buys and shorts the bottom $\tilde{\alpha}$ quintile's strong buys (black solid line) or (2) buy the top $\tilde{\alpha}$ quintile's strong sells and shorts the bottom $\tilde{\alpha}$ quintile's strong sells (red dashed line), or (3) buys the first hedge portfolio of strong buys and shorts the second hedge portfolio of strong sells (blue dashed-dotted line).

(a) Strong Buys - Strong Sells

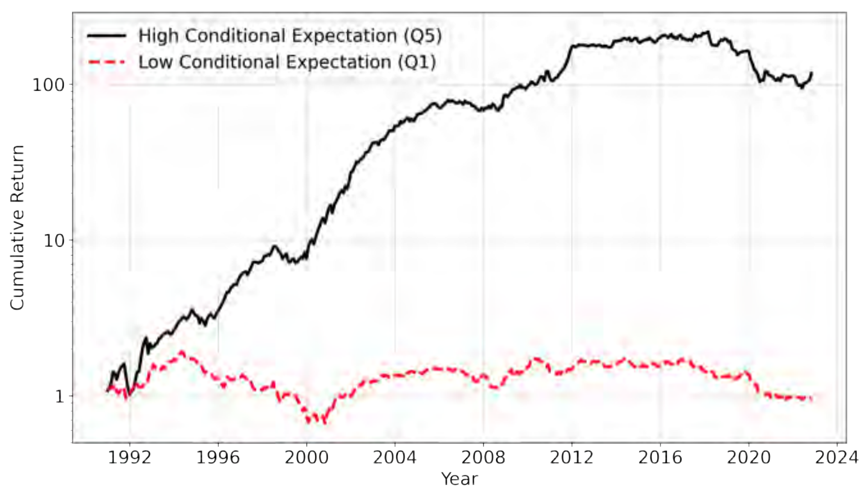

(b) High Minus Low Conditional Expectation

Figure 7: Cumulative Returns (Daily)
This figure plots cumulative returns for portfolios formed by sorting on (1) an insider's conditional expectation $\tilde{\alpha}$ and (2) whether the trade is a buy or a sell. The mixture model is estimated in an expanding window fashion each year using the latest average abnormal return and standard error for each insider with at least ten trades prior to that year's end. Based on the estimated parameters and each insider's average abnormal return and standard error, the conditional expectation of insider informed trade is calculated. Insiders are sorted into quintiles on the basis of the conditional expectation. This sorting is done based on an insider's trade history up to the prior year's end for a given date's portfolio formation. The second sort is based on whether the trade is a buy or a sell. A stock from a given trade enters the buy or sell portfolio the day following the trade and is included for forty days. The portfolio holds all stocks included in the portfolio that day at equal weights. Panel A reports cumulative returns for hedge portfolios that buy stocks with strong inside buying pressure and sell stocks with strong inside selling pressure. The black solid (red dashed) line represents this strategy for insiders in the top (bottom) quintile of ex-ante conditional expectation. Panel B reports the cumulative performance for hedge portfolios that either (1) buy the top $\tilde{\alpha}$ quintile's buys and shorts the bottom $\tilde{\alpha}$ quintile's buys (black solid line) or (2) buy the top $\tilde{\alpha}$ quintile's sells and shorts the bottom $\tilde{\alpha}$ quintile's sells (red dashed line), or (3) buys the first hedge portfolio of buys and shorts the second hedge portfolio of sells (blue dashed-dotted line).

(a) Buys - Sells

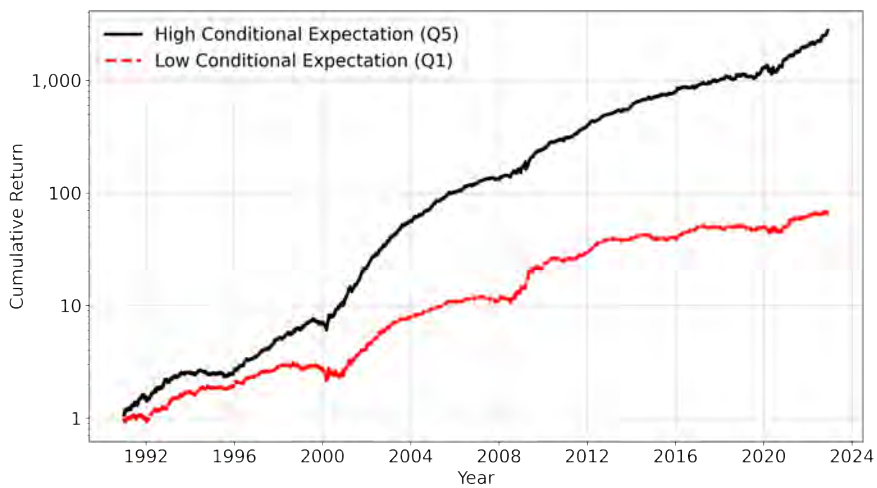

(b) High Minus Low Conditional Expectation

Figure 8: Insiders' Information and Convergence to Market Efficiency

This figure plots cumulative returns for portfolios formed by sorting on (1) an insider's conditional expectation $\tilde{\alpha}$ and (2) whether the trade is a buy or a sell, as in Figure 7. Within each panel, portfolio formation differs only in the entry date of a stock into the portfolio. A stock from a given trade enters a portfolio either 1 (black solid line), 3 (blue dashed), or 5 (red dashed-dotted line) trading days following the trade date; in each case, the stock leaves the portfolio forty trading days after the trade date. Panel A reports cumulative returns for the hedge portfolio that buys stocks with strong inside buying pressure and sells stocks with strong inside selling pressure for insiders in the top quintile of ex-ante conditional expectation. Panel B reports the cumulative performance for the hedge portfolio that buys the top $\tilde{\alpha}$ quintile's buys and shorts the bottom $\tilde{\alpha}$ quintile's buys.

(a) High Conditional Expectation (Q5): Buys - Sells

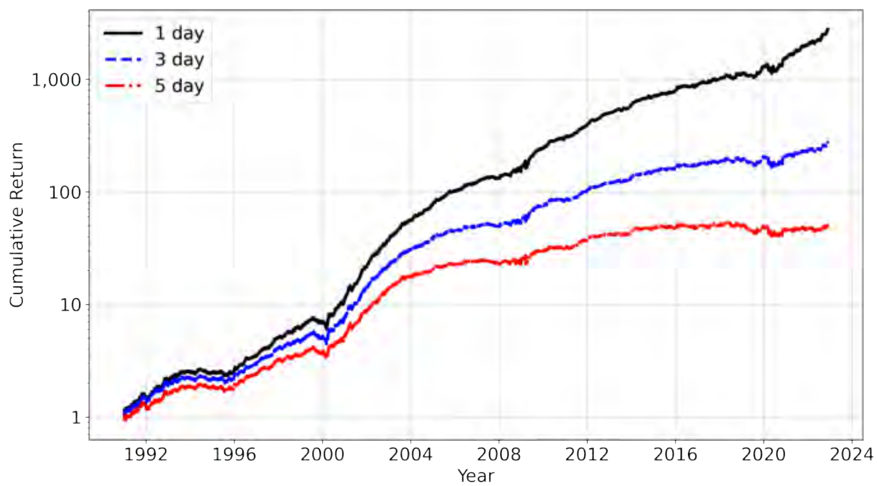

(b) High Minus Low Conditional Expectation: Buys

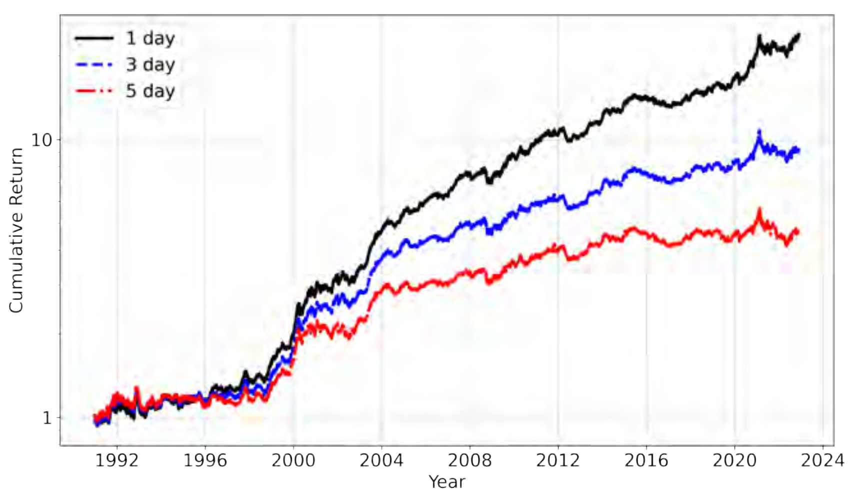

Figure 9: Jumps in Information Acquisition by Sophisticated Market Participants
The figure reports the weekly number of insider trading disclosures accessed by the indicated financial institutions.

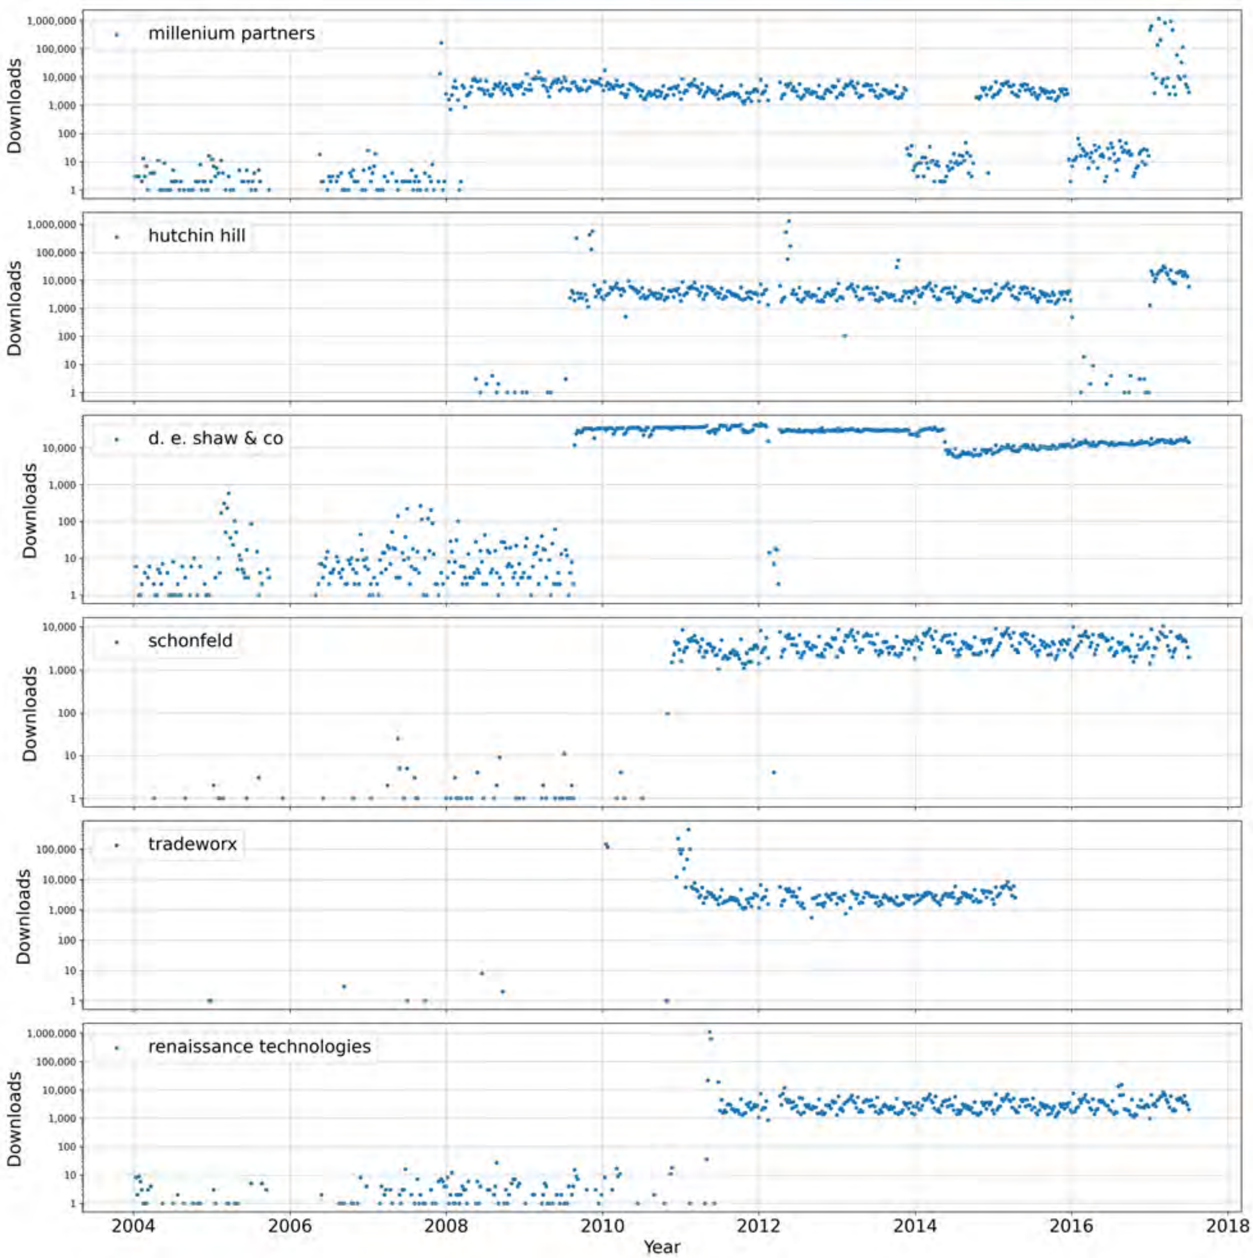

Figure 10: Conditional Probability an Insider's Trade is Informed

The figure plots the conditional probability that a trade made by a corporate insider was informed as a function of the realized trade return r_{ij} and attributes of the insider's past trading history. Specifically, the probability is conditioned on the insider's past average abnormal return $\bar{r}_i$, the standard deviation of their previous returns σ_i, and the number of past trades. Each panel shows the conditional probability curves for the number of trades and standard deviation in the panel header and for past average returns of -1%, 0%, and 1%.

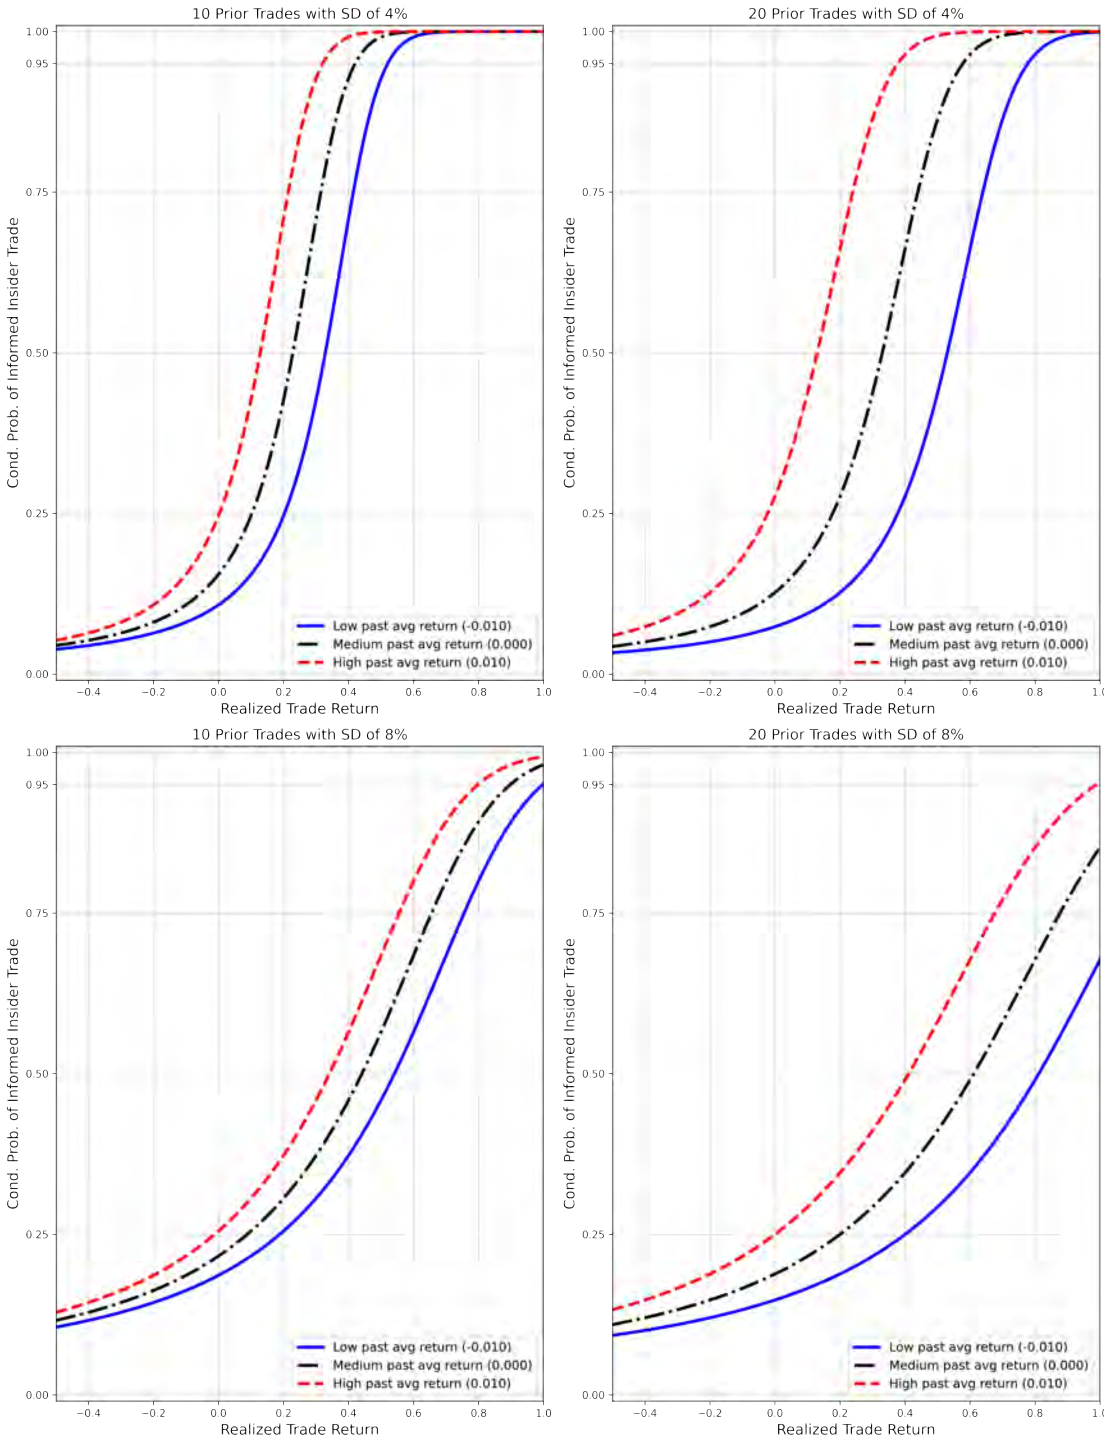

Figure 13: Conditional Probability an Insider Trade is Informed
The figure plots the trade-level conditional probability (15) that a trade was informed. The sample contains all trades made by corporate insiders from January 1990 through November 2022.

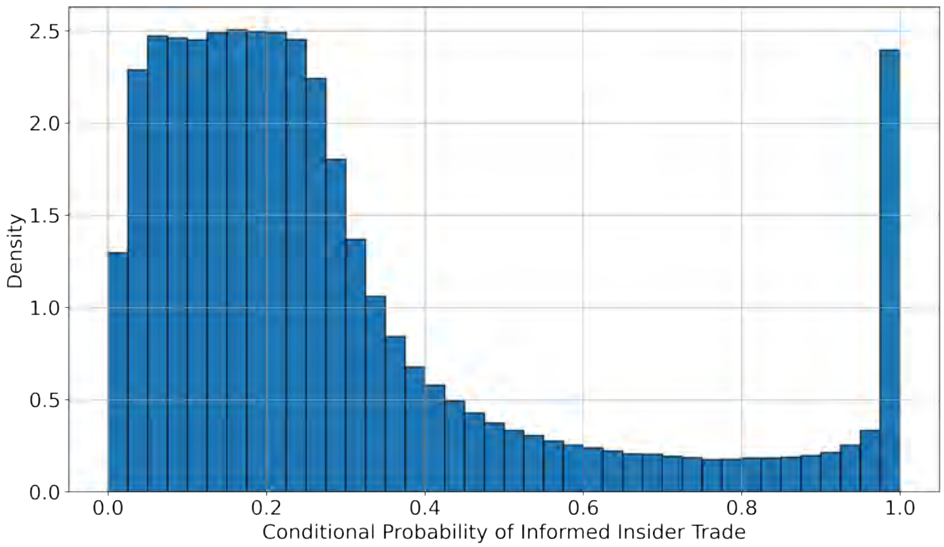

Figure 15: Case Study: Enron

The figure plots conditional probability (top figure in each panel) and the trade-by-trade evolution of the insider's past average abnormal return and 95% confidence band (bottom figure in each panel). Panels (a) and (b) report these time-series for Enron executives Kenneth Lay and Jeffrey Skilling, respectively.

(a) Kenneth Lay

(b) Jeffrey Skilling