Detecting Informed Trade by Corporate Insiders

Abstract

We introduce a mixture model that leverages the cross-section of insiders' trade histories to infer which insiders are more likely to engage in informed trade. The estimation explicitly accounts for the profitability and noisiness of insiders' past performance. Out-of-sample returns are higher for stocks traded by insiders identified as more likely to use information. Prices reflect this information faster over the last decade. The model for insiders informs a person-specific mixture distribution that is used to classify whether *any* disclosed trade is informed. Whether trades are prescheduled, option-related, or by inside blockholders significantly relates to the probability they are informed.

March 17, 2025

1. Introduction

Corporate insiders—such as managers and directors—possess private information about their firms, giving them a potential advantage in financial markets. To mitigate the risks associated with insider trading, which can reduce market liquidity and deter investor participation, the Securities and Exchange Act of 1934 introduced Rule 10b-5. However, enforcement remains challenging and costly, as insiders trade for various reasons, many of which are entirely legal. For instance, insiders often hold significant portions of their wealth in company stock and may trade for liquidity or diversification rather than due to an informational advantage. Compounding this difficulty, stock returns are inherently noisy. Some informed trades may appear unprofitable by chance, while profitable trades may just be lucky rather than informed.

In disentangling information from noise, regulators, traders, and academics all face a significant challenge: whether trade by corporate insiders is informed is not directly observable. However, 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 potential 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 profitability and noise differ across insiders. We take advantage of the fact that there is additional information in the cross-section of signals to identify the unconditional distribution of informed insider trading. 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 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).

[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), and security analysts (Crane and Crotty, 2020).

Using the estimated mixture model parameters and the 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 an insider's estimated probability of informed trade moves off the unconditional estimate as a function of the magnitude and precision of the realized returns. Intuitively, the econometrician should update more strongly that an insider trades on private information if the insider's average return is higher. However, 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 trade profitability expectations predict differences in stock returns. After sorting on this conditional expectation, 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, 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.

Identifying informed trade is complicated by the fact that competing traders may try to profit off insiders' trading behavior. We show that prices have become more efficient over time with respect to the trading behavior of informed insiders. 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. One could easily mistake this change in horizon for a reduction in the amount of informed trading. 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 return predictability of the mixture model's 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 that utilizes an insider's probability of trading on information based on the insider-level mixture model. 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 *all* trades, including those made 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

[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.

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, we find that pre-scheduled 10b5-1 purchases are more likely to be informed, but pre-scheduled 10b5-1 sales and option-related executions are less likely to be informed. Further, CEOs, CFOs, and large inside blockholders make more informed trades. Through a case study of Enron, we show how different return histories for insider 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 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

[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).

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., 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, a 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, so 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 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 all 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 ($\alpha_i = 0$) and that the true average abnormal return of informed insiders is distributed exponentially with

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

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 illustrates 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.

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 $f_{\text{I}}(\bar{r}_i|\text{informed}, s_i)$ 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:

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

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

[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.

[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.

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

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 $\bar{r}_i$, their standard error s_i, and the estimated parameters. Denote the conditional probability by $\tilde{\pi}_i$. The conditional probability that insider i makes informed trades is:

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

Let $\tilde{\mu}_i$ denote the conditional expectation of insider i's profitability *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 Internet Appendix A.1 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 from uninformed insiders, while positive abnormal returns are more likely to be from 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. A naive proxy one may think of using instead would be to calculate the t-statistics from $\bar{r}_i$ and s_i or to just use $\bar{r}_i$. However, as Figure 2 makes clear, due to the noise in returns, especially when the noise is high, you could estimate negative average returns and t-statistics for insiders making informed trades due to bad luck. Our methodology correctly accounts for this possibility, which we show makes a meaningful difference when predicting returns out-of-sample in Section 3.

The effects of estimation noise on 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 to the unconditional estimate 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, 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.

Table 1 reports distributional statistics of the average abnormal return $\bar{r}_i$ and the standard error s_i. 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. 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.

[7] In our full-sample estimation, n_i is simply the total number of distinct stock-date observations for insider i. Our out-of-sample estimation estimates 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.

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, $\hat{\pi}$, is 28.6%. The average abnormal profitability by insiders who trade on information is, $\hat{\mu}$, is 3.6% per month.

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 parameters from the unconditional distribution, π and μ.

Figure 3 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 3 shows that an econometrician (or regulator) should make very different inferences 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 3 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 abnormal profitability 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 towards the unconditional estimate because the conditional probability that the insider uses information is lower due to the higher likelihood of luck.

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 to that year's end. The time-series of π and μ are plotted in Figure B.1 of the Internet Appendix. 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 profitability is more variable, dropping from above 5% in the late 1980s 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.

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.

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 that given stock. For instance, Buy Quintile 5 (Sell Quintile 5) 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 3 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 a 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 a cross-sectional spread in future returns as a function of an insider's 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)). The results are also robust to using portfolio sorts instead of regression analysis. This additional analysis is reported and discussed in Internet Appendix B.1. Overall, the results of Table 3 provide strong support for the mixture model's ability to differentiate between insiders with higher propensities to profit from their private information.

It is natural to ask how the mixture model approach compares to alternative methods that account for estimation noise differently. In the extreme, the econometrician could ignore noise in the past average return and simply sort insiders based on their past average profitability, $\bar{r}_i$. Panel A of Table 4 reports the Hi-Lo spreads when the regression analysis is performed using this sorting algorithm. There is no predictive spread between Hi-Lo insiders; any information in past average returns alone is swamped by the noisiness of return histories. One could instead scale average returns by their standard error, sorting insiders by the t-statistic of their past profitability, $\bar{r}_i/s_i$. Panel B of Table 4 shows that such an approach leads to point estimates in the correct direction, but the classification is still too noisy, as evidenced by Hi-Lo p-values far from significance. By comparison, the mixture model's use of information from the cross-section of insider profitability and noise substantially improves the predictive power of insider's trades (Panel C). t-statistics shrink towards a null of zero profitability, whereas our methodology shrinks towards the unconditional mixture distribution estimated from the cross-section of insiders. Therefore, shrinkage and using the information recovered from the cross-section of insiders is important in efficiently measuring informed trade.

3.1. 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 4 shows cumulative returns from several monthly portfolio strategies. We first sort by an insider's conditional expectation $\tilde{\alpha}_i$. In DeMarzo et al. (1998), an optimal enforcement regime only investigates after a large trading volumes (or large price movements or both). Following this intuition, we additionally sort trades each month into quintiles based on a trade's signed trading strength, defined as the signed trade volume divided by the overall trading volume in the stock that month. The top panel of Figure 4 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 indicates 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 are 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 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, 25% of insiders are estimated to make informed trades. These estimates are plotted in Figure B.2 of the Internet Appendix.

What can then 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. Consistent with this, we find 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. To gauge how quickly, we consider portfolios formed daily, waiting either 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 5. 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. This is evidence that the information in these trades is getting into prices much more quickly.[8] Interestingly, this coincides with dramatically increased information acquisition by sophisticated financial institutions around the same time the monthly strategy performance declines, as evidenced by analysis of the SEC EDGAR server logs. These additional analyses are reported and discussed in Internet Appendix B.2.

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 probability and expectation measures 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; insiders with some buying and selling activity within the year are classified as either medium- or short-term insiders. High quarterly earnings announcement (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.

[8]Until 2002, insiders had 10 days to file a Form 4 after a trade. Part of the Sarbanes-Oxley Act in 2002 reduced the reporting window to 2 days.

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's conditional probability $\tilde{\pi}_i$ estimates correlate with these measures. The results are reported in Panel A of Table 5. Consistent with the conclusions of the prior literature, we find positive correlations between each proxy and the conditional probability $\tilde{\pi}_i$. Non-routine insiders have 4% higher $\tilde{\pi}_i$, 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 $\tilde{\pi}_i$, and short-horizon insiders have 9% higher $\tilde{\pi}_i$, compared to long-horizon insiders. Finally, insiders in the top quintile of QEA profitability exhibit $\tilde{\pi}_i$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 ($\tilde{\alpha}_i$) as well (Panel B of Table 5). 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 $\tilde{\alpha}_i$ is 29 bps higher than the 60 bps average $\tilde{\alpha}_i$ of routine insiders. The wedge is even larger for long- versus short-horizon insiders. Long-horizon insiders have an average $\tilde{\alpha}_i$ of 57 bps; short-horizon insiders' average $\tilde{\alpha}_i$ 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 $\tilde{\pi}_i$ and $\tilde{\alpha}_i$. This suggests that our measure is capturing different information than prior work. In fact, when we include these proxies as controls in our out-of-sample predictability analysis, they have little effect on the predictability of the mixture model. On the other hand, inclusion of trading indicators from 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. These results are tabulated in Internet Appendix Table B.3 and discussed in detail in Internet Appendix B.4.

4.2. 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 6 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.[9] 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 considered in Section 5.4.

5. Detecting Which Trades Are Informed

5.1. A Trade-Level Mixture Model

In this section, we are interested in classifying whether any *individual* trade made by an insider was potentially informed or not. To do so, we estimate a probabilistic model that leverages each insider's return history and the cross-section of returns up to that point. Specifically, we use the expanding window estimates of the insider-level mixture distributions π and μ, and each insider i's average return $\bar{r}_i$ and standard error s_i resulting from the insider's $j-1$ previous trades. We

[9] It is worth noting that a substantial fraction of traders are unclassified by prior literature. Given the amount of unclassified traders who do sometimes trade off information, it is important to consider these traders when studying insider trading.

model the individual return, r_{ij}, for the j^{th} trade by the i^{th} insider as:

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

We assume that ε_{ij} is normally distributed with mean zero and standard deviation σ_i. Label the standard deviation of insider i's return history, $\sqrt{(j-1)} \cdot s_i$, as $\tilde{\sigma}_i$. For insiders with a reasonably large history ($j > 10$), we set the trade-level standard deviation σ_i equal to $\tilde{\sigma}_i$. 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 $\tilde{\sigma}_i$ and the cross-sectional median standard deviation, $\text{med}(\tilde{\sigma})$, 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 $\text{med}(\tilde{\sigma})$.

We use the history of trades to inform the probability that α_{ij} is positive. Denote the probability that $\alpha_{ij} > 0$ as p_{ij}. We allow this probability to vary depending on whether the trade is a purchase or a sale. For both trade types, we also let p_{ij} vary as a function of the insider's conditional probability of using information $\tilde{\pi}_i$. This allows the insider's past history of trades to inform the model. For insiders trading for the first time, we use the estimated π, the unconditional estimate, as the insider's $\tilde{\pi}_i$ estimate, which implicitly utilizes information from the history of the cross-section of past traders. The probability that $\alpha_{ij} > 0$ is:

$$p_{ij} = \frac{\exp(b_0 + b_1\tilde{\pi}_i + b_2\mathbf{1}_{buy} + b_3\tilde{\pi}_i\mathbf{1}_{buy})}{1 + \exp(b_0 + b_1\tilde{\pi}_i + b_2\mathbf{1}_{buy} + b_3\tilde{\pi}_i\mathbf{1}_{buy})} \,. \tag{12}$$

The coefficients b_0, b_1, b_2, and b_3 are parameters that we estimate via maximum likelihood. 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 also helps to inform the distribution of α_{ij} conditional on a positive α_{ij} realization. For insiders trading for the first time, the conditional distribution is simply the unconditional distribution of informed trade; that is, exponential with mean μ/p_{ij}.[10] 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 $\bar{r}_i - \sigma_i^2/((j-1)\cdot\mu/p_{ij})$ and standard deviation $\sigma_i/\sqrt{j-1}$.

[10]Note that we have scaled the conditional mean parameter μ by p_{ij}. Recall that the estimate of μ captures the *average* abnormal return across trades. In the trade-level model, this average can be approximated as the product of the probability a given trade is informed (p_{ij}) and a trade-level conditional mean. Thus, we use a trade-level conditional mean of μ/p_{ij}.

The density of r_{ij} conditional on a positive α_{ij} realization is the convolution of α_{ij} and ε_{ij}:

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

In Internet Appendix A.2, we show that the convolution of a truncated normal and normal distribution for $h_{\mathrm{I}}(r_{ij}|\alpha_{ij} > 0, \sigma_i, p_{ij})$ with $j > 1$ can be expressed in closed form, which substantially eases the computational burden of estimating the trade-level model.

The likelihood of observing return r_{ij} is a mixture:

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

We estimate the trade-level model via maximum likelihood using a sample of insider trades from 1991 through 2022. The estimated coefficients in Equation 12 are $\hat{b}_0 = -3.1$, $\hat{b}_1 = 4.1$, $\hat{b}_2 = 1.4$, $\hat{b}_3 = -1.9$. For both buys and sells, the estimated trade-level probability is increasing in the ex-ante probability $\tilde{\pi}_i$ that an insider uses information. The estimated trade-level probability p_{ij} is higher for buys than for sells for most empirically observed levels of $\tilde{\pi}_i$. The estimated functional form of p_{ij} is displayed in Figure B.5 of the Internet Appendix.

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

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

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 $\tau(r_{ij}|\sigma_i, p_{ij}) > 0.5$ are more likely informed than uninformed. Regulators with a limited investigative and enforcement budget might choose a higher threshold when considering which trades to investigate further.

5.2. Conditional Probabilities and Return Thresholds

An important feature of this model for insider trades is that the the conditional probability a trade is informed 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 and return standard deviation. Figure 6 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 $\bar{r}_i$ and the standard deviation of their past trades σ_i.

As is natural, higher realized returns translate into higher conditional probabilities that a given trade was informed. Each panel fixes the prior return standard deviation and whether the trade was a purchase or a sale. Consider the top left panel, which considers a purchase made by an insider with a history of 10 trades with a standard deviation of 6%. 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 25% for the model to classify the trade as likely informed ($\tau_{ij} > 0.5$). This threshold is higher (lower) for insiders who have lost (gained) 2% on average on their past 10 trades. For sales, these differences increase, and even higher returns are needed to classify the trade as informed (compare the top and bottom rows). This is because the estimated ex-ante probability p_{ij} that a sale is informed is lower than that for a purchase ($\hat{b}_2 > 0$).

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 left columns to the right column of Figure 6). 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 7, which plots the trade-level return thresholds at which the current trade is classified as likely informed ($\tau_{ij} > 0.5$). This corresponds to a preponderance-of-evidence burden of proof. Each panel plots the thresholds as a function of an insider's historical average return ($\bar{r}_i$) and of the standard deviation of the insider's past trade returns, σ_i. The low, median, and high σ_i levels correspond to the 10th, 50th, and 90th percentiles of empirical trade-level standard deviations for a given (binned) past average return. The top (bottom) panel shows thresholds for purchases (sales).

Empirically, σ_i is strongly positively correlated with the absolute value of the past average returns; more extreme $\bar{r}_i$'s are associated with noisier underlying returns. As described above, the return threshold generally declines in past average return but increases in past return standard deviation. For insiders with negative past average returns, these two effects work in the same direction, and the return threshold declines as past average returns become less negative. For insiders with positive past average returns, however, the two effects work in opposite directions. As a result, the threshold for median σ_i still declines with increasing $\bar{r}_i$, but it does so at a lower rate than for insiders with negative $\bar{r}_i$. For insiders with high levels of σ_i, the noise effect can dominate the past average return effect, leading to an increase in the threshold as $\bar{r}_i$ increases.

One thing to note from these figures is that for some histories, the return threshold can be very high. This is particularly true for insider sales. This is because 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 median return standard deviation, a current purchase would need an abnormal return of around 30% over the next month to be classified as likely informed. A sale made by the same insider would need an almost 50% abnormal return over the next month to be classified as likely informed. 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_{ij} from equation (15) for a sample of *all* insider trades from 1991 through 2022. The probability depends on the estimates of b_0, b_1, b_2, and b_3 (reported in Figure B.5) and on population parameters π and μ estimated using expanding windows using all data up until the prior year-end (plotted in Figure B.1).

Figure B.6 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 7 reports averages of the conditional probabilities τ and as well as an indicator for whether a trade is classified as likely informed ($\tau > 0.5$). The average conditional probability is about 20%, and about 10% of trades are likely informed. 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 (30%) than for sales (15%), and about 15% of buys clear the 0.50 threshold compared to only 8% of sales.

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 the 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.

5.4. Applications of Trade-level Conditional Probabilities

The trade-level conditional probabilities allow researchers to explore economic questions of interest related to corporate insider trading. Table 8 reports regression results of likely informed trades for buys or sells on characteristics of the trade and the insider that are likely to relate to information content of the trade.[11]

Pre-scheduled 10b5-1 trading plans are one potential policy remedy to mitigate informed trading by corporate insiders. These plans are not without controversy, however, as insiders may strategically cancel planned trades (Jagolinzer, 2009; Lenkey, 2019). Our results suggest that purchases made under the 10b5-1 plans are more informed than other purchases, while 10b5-1 sales are less informed than other sales. These results hold even with insider-level fixed effects.

Results are also heterogeneous across buys and sales with respect to the magnitude of the trade in dollars. Larger-sized purchases are more likely to be informed, while larger sales are less likely to be informed. This is consistent with markets interpreting larger sales as more likely to be liquidity motivated. For sales, we are also able to control for whether or not the selling is related to an option position. Option-related sales are less likely to be informed, consistent with insiders diversifying following option vesting. Most of this relation is due to cross-sectional differences across insiders, as the point estimate drops substantially when insider fixed effects are included.

We also consider whether sustained trading by an insider or correlated trading by others at the same firm is related to elevated adverse selection due to insiders. For both buys and sells, sustained trading by an insider is positively related to an elevated probability of an informed insider trade. There does not seem to be a strong relation between other insiders trading and the probability of an informed trade. Buys are slightly more likely to be informed if others at the firm are also buying while sells are less likely to be informed if others at the firm are also selling. This suggests clustered selling (buying) is likely liquidity-driven (information-driven).

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

[11] Internet Appendix Table B.4 replaces the dependent variable with the conditional probability of an informed trade.

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. While prior work documents that the profitability of some types of insiders' trades is higher (e.g., CFOs' Wang et al., 2012), it is still unclear how this relates to the prevalence of informed trade at the transaction level. That is, is the CFOs' outperformance due to consistent information-based trading, or is it due to occasional valuable trades mixed in among liquidity-driven transactions? While this distinction is key to understanding the adverse selection issues inherent in markets, it is challenging to address without a trade-level measure of informed trading. Using the roles in the firm as disclosed on Form 4s, we test if likely informed trades are a function of their role. The results for buys and sells are tabulated in Panel B of Table 8.

Both buys and sells made by CEOs and CFOs are more likely to be informed. Buys made by chairmen are also more likely to be informed buys. Interestingly, the insider characteristic most associated with informed trade is whether or not the insider holds a block of at least 10% of shares. The probability a purchase (sale) is likely informed is 13% (9%) higher for these insiders.

5.5. A Case Study

We now consider how the model classifies trades made by two prominent insiders accused of informed insider trading: 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 of nine other insider trading charges.

For each Enron executive, Figure 8 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, Enron's share price 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 does not classify any of Skilling's later trades as likely to have been informed, while almost all of Lay's trades are classified as likely informed over the same period.

6. Conclusion

Corporate insiders have access to private information about their firms. We employ mixture model methods that leverage the cross-section of insiders and their trading histories to classify which insiders are likely trading on information and which trades are most likely to contain information. Importantly, these methods account for the significant amount of noise in trading performance. Out-of-sample tests confirm that our measure provides new economic content regarding the information embedded in trades by insiders identified as likely to trade on private information, even after controlling for prior literature's proxies.

In portfolio tests, we show that the market has incorporated this information much more quickly since 2012. This shift is likely due to the rise of scraping Form 4 filings by major hedge funds beginning around that time. When we augment our model with proxies from prior literature, we confirm that insiders classified as informed are indeed more likely to trade on private information. However, insiders excluded due to sample selection criteria appear just as likely to trade on information as those classified as informed.

Our insider-level model generates an insider-specific mixture model for trade returns, helping determine whether a trade was likely informed. We find that sales related to 10b5-1 plans and option-related sales are less likely to be informed. Conversely, trades made by CEOs, CFOs, and inside blockholders are more likely to contain private information. Finally, the model's conditional probability of an informed trade can help regulators allocate enforcement resources more effectively when monitoring corporate insider trading.

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.

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.

Jagolinzer, Alan D, 2009, Sec rule 10b5-1 and insiders' strategic trade, *Management Science* 55, 224–239.

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.

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

Lenkey, Stephen L, 2019, Cancellable insider trading plans: an analysis of sec rule 10b5-1, *The Review of Financial Studies* 32, 4947–4996.

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

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	54,274		

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: 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 up to that month, 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	180,715	180,715	180,715	180,715	180,715	180,715
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 4: Return Predictions Using Average Return, t-statistic, and Conditional Expectation

This table reports comparisons of regressions of monthly stock returns as a function of insider buying and selling activity in the prior month. Insiders are sorted into quintiles each month based on their past average return (Panel A), the t-statistic of their past average return (Panel B), or their conditional expectation (Panel C), as in Table 3. Each measure is estimated in an expanding window fashion each year, as described in Table 3. The regression specification corresponds to the last column of Table 3. The table reports the regression coefficients on the High and Low quintile indicators for buys and sells, their difference, and the p values for tests of whether the coefficients on the High and Low quintiles differ. Panel C of this table repeats information from the last column of Table 3 for ease of comparison.

Panel A. Average Return

	Low (Q1)	High (Q5)	Hi-Lo	p-value
Buy	0.89	0.90	0.01	0.7955
Sell	-0.24	-0.23	0.01	0.9615

Panel B. t-statistic

	Low (Q1)	High (Q5)	Hi-Lo	p-value
Buy	0.48	0.62	0.14	0.5456
Sell	-0.09	-0.22	-0.13	0.3035

Panel C. Mixture Model Conditional Expectation $\tilde{\alpha}$

	Low (Q1)	High (Q5)	Hi-Lo	p-value
Buy	0.11	0.79	0.68	0.0059
Sell	0.07	-0.23	-0.30	0.0381

Table 5: 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	130,671	111,582	24,639

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	130,671	111,582	24,639

Table 6: 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 7: 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_{ij}. The second column reports the sample average of an indicator for whether the conditional probability is greater than 50% (Likely Informed). The statistics are reported separately for the overall sample, for purchases and sales, and for quintiles of trading strength. Trading strength is the signed trading volume divided by the stock's monthly volume. Trading strength quintiles are formed monthly.

	Average Conditional Probability	Fraction Likely Informed
Overall	0.1993	0.1021
Trade Direction		
Purchases	0.2958	0.1497
Sales	0.1542	0.0795
Trade Strength		
Strong Sell	0.1771	0.1080
2	0.1565	0.0796
3	0.1584	0.0724
4	0.2074	0.0905
Strong Buy	0.2974	0.1614

Table 8: Likely Informed Trades, Trade Characteristics, and Insider Roles

This table reports regressions of an indicator of likely informed trades ($\mathbb{1}(\tau_{ij} > 50\%)$) as defined Section 5) on trade level characteristics in Panel A and their role(s) within the firm in Panel B. In Panel A, our independent variables include if the trade was a scheduled 10b5-1 trade, which is voluntarily disclosed, the log of the trade size in dollars, if the sale was related to an option execution, and the number of other traders in the firm or days the insider traded in the past five days. In Panel B, we include a dummy variable for roles disclosed by the insider or if they were a direct owner in the firm. We control for prior proxies of informed insiders and denote which fixed effects are used at the bottom of each row. 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: Trade Characteristics

	Buy $\mathbb{1}(\tau_{ij} > 50\%)$		Sell $\mathbb{1}(\tau_{ij} > 50\%)$	
	(1)	(2)	(3)	(4)
10b5-1 Trade	0.0613	0.0487**	-0.0176***	-0.0048**
	(1.44)	(2.30)	(-3.55)	(-2.21)
log(Trade Size)	0.0100***	0.0042***	-0.0062***	-0.0042***
	(6.65)	(4.49)	(-7.85)	(-10.41)
Option Related Sell			-0.0245***	-0.0016*
			(-10.00)	(-1.90)
# Days Buying in Past Week	0.0486***	0.0099***		
	(8.82)	(6.22)		
# Other Traders Buying in Past Week	0.0023	0.0037**		
	(1.40)	(2.70)		
# Days Selling in Past Week			0.0343***	0.0114***
			(17.42)	(17.38)
# Other Traders Selling in Past Week			-0.0024**	-0.0013***
			(-2.31)	(-4.35)
Controls for Prior Proxies	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Insider FE	N	Y	N	Y
Adj. R^2	0.0857	0.4366	0.0475	0.4902
Observations	667,570	624,894	1,408,953	1,374,822

Panel B: Insider's Role

	Buy $\mathbb{1}(\tau_{ij} > 50\%)$	Sell $\mathbb{1}(\tau_{ij} > 50\%)$
	(1)	(2)
CEO	0.0454*	0.0245***
	(2.72)	(3.70)
CFO	0.0150*	0.0101*
	(2.12)	(2.68)
Inside Block > 10%	0.127***	0.0944***
	(6.90)	(9.95)
Chairman	0.0452*	0.00133
	(2.17)	(0.22)
Director	-0.0279*	-0.00486
	(-2.35)	(-1.25)
Officer	-0.0205*	-0.0274***
	(-2.26)	(-7.80)
Officer and Director	-0.0188	-0.00955
	(-1.43)	(-1.93)
Vice Presidents	0.00350	-0.00837**
	(0.43)	(-3.40)
Direct Ownership	-0.0321*	-0.0275***
	(-2.53)	(-6.97)
Controls for Prior Proxies	Y	Y
Year FE	Y	Y
Adj R^2	0.0678	0.0325
Observations	668,912	1,411,800

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

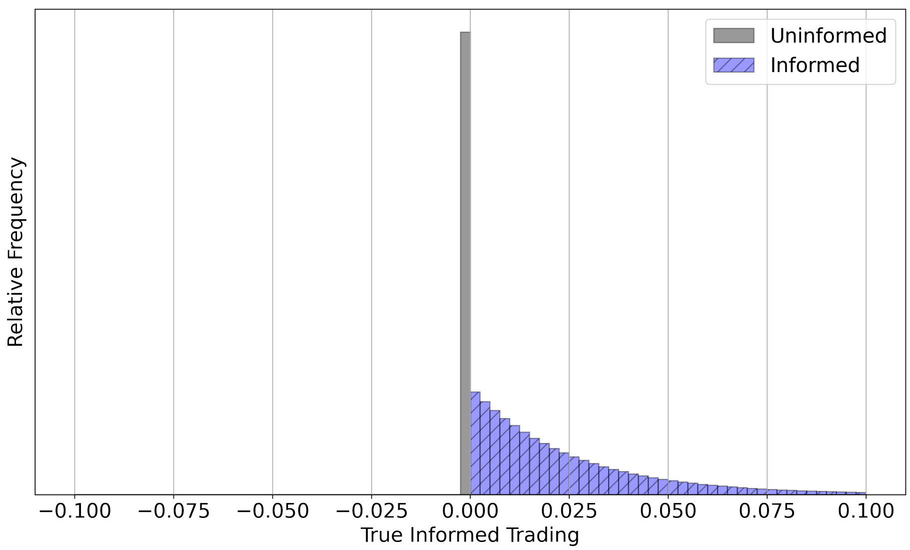

(b) Estimated Abnormal Return

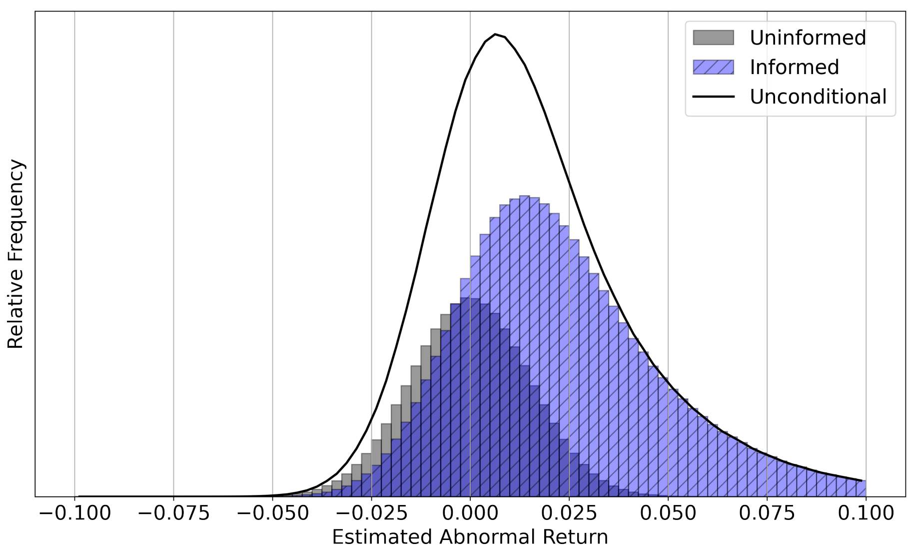

Figure 2: Conditional Probabilities and Expectations

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

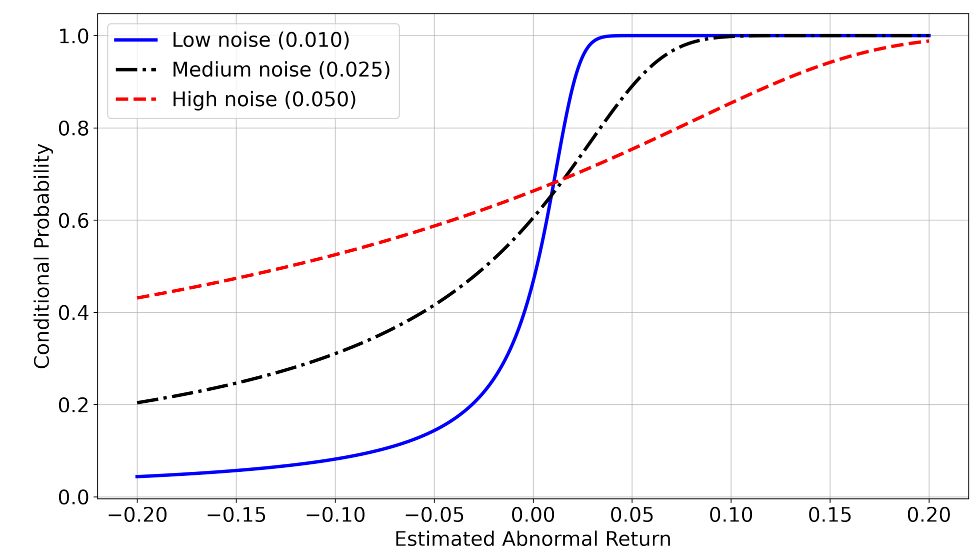

(b) Conditional Expectation of Insider Information

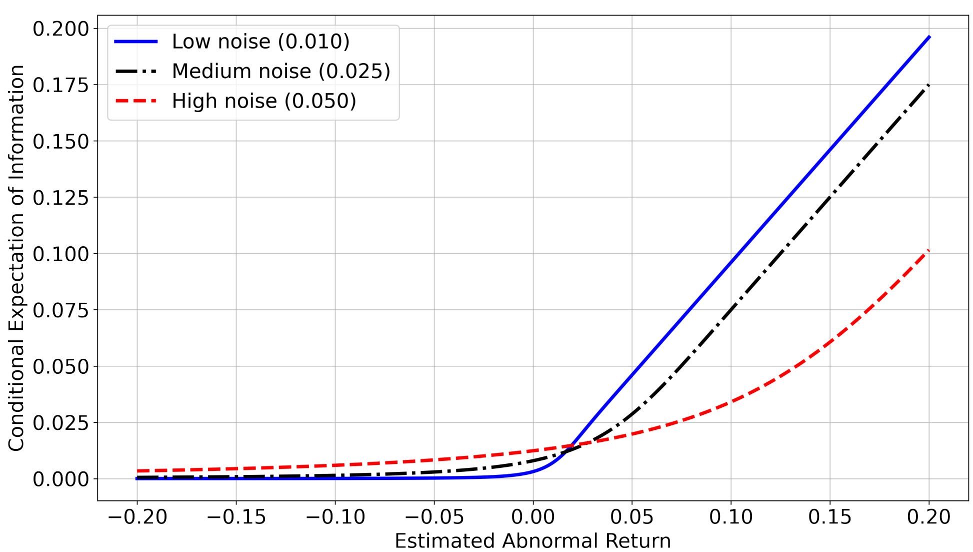

Figure 3: Noise and Empirical Conditional Probabilities and Expectations

This figure illustrates conditional probabilities and expectations as a function of the estimated average abnormal return and its standard error (i.e. its noise) using the estimated mixture model parameters from the full sample of insiders with at least 10 trades. The standard errors used for each line correspond to the indicated values and percentiles in the sample. 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.

(a) Conditional Probability Insider is Informed

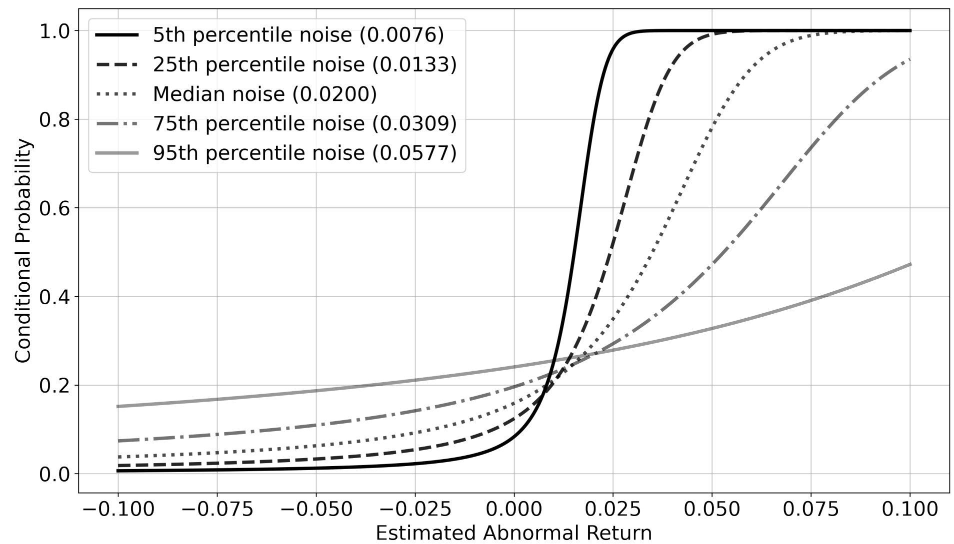

(b) Conditional Expectation of Insider Information

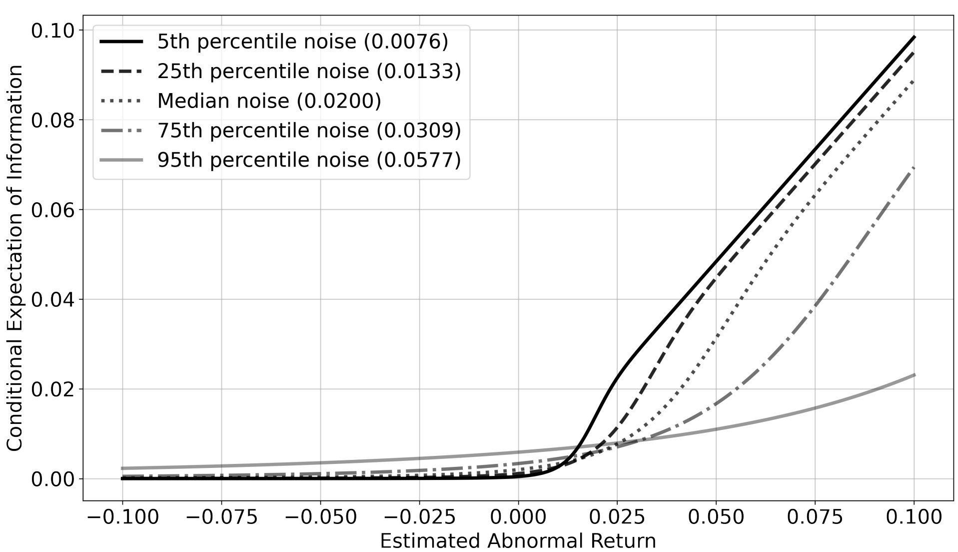

Figure 4: Cumulative Returns (Monthly)

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. 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 $\tilde{\alpha}_i$ quintile's strong buys and shorts the bottom $\tilde{\alpha}_i$ quintile's strong buys (black solid line) or (2) buy the top $\tilde{\alpha}_i$ quintile's strong sells and shorts the bottom $\tilde{\alpha}_i$ 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

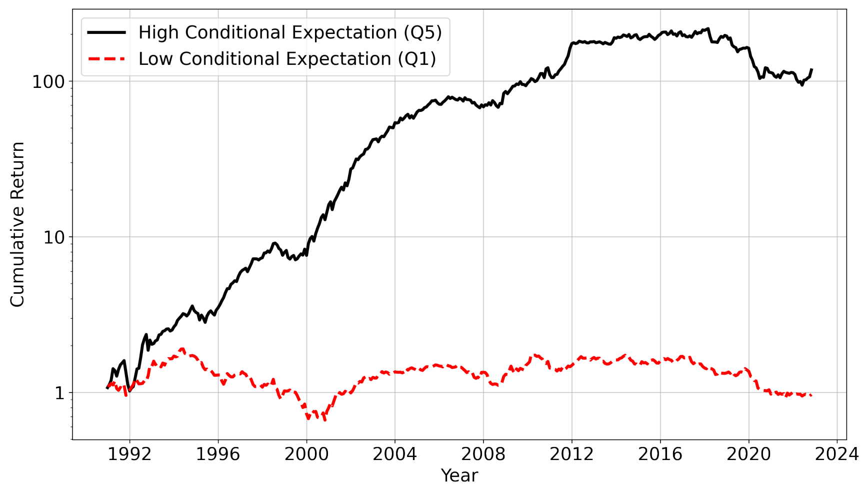

(b) High Minus Low Conditional Expectation

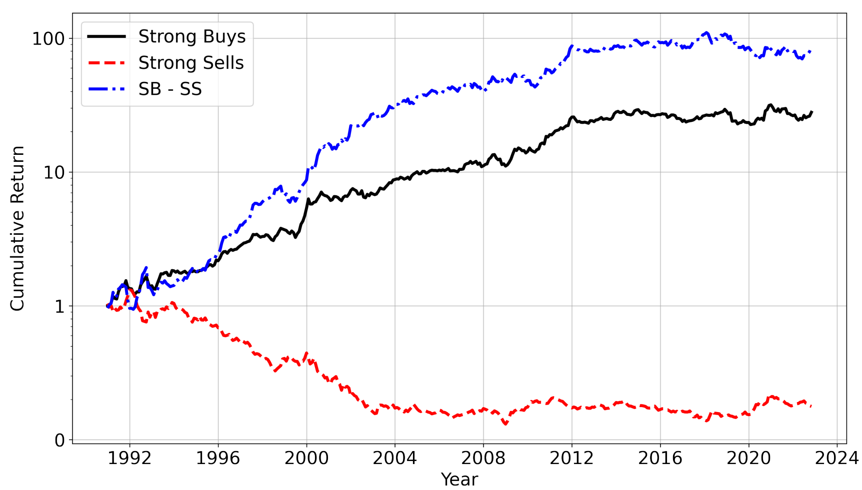

Figure 5: Insiders' Information and Convergence to Market Efficiency

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. A stock from a given trade enters the buy or sell portfolio at the indicated number of days following the trade. The portfolio holds all stocks included in the portfolio that day at equal weights. 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 $\tilde{\alpha}_i$ quintile's buys and shorts the bottom $\tilde{\alpha}_i$ quintile's buys.

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

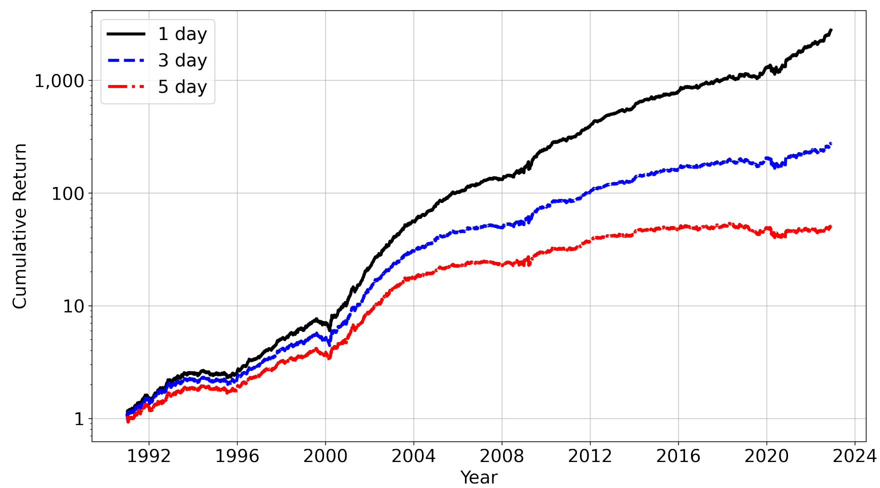

(b) High Minus Low Conditional Expectation: Buys

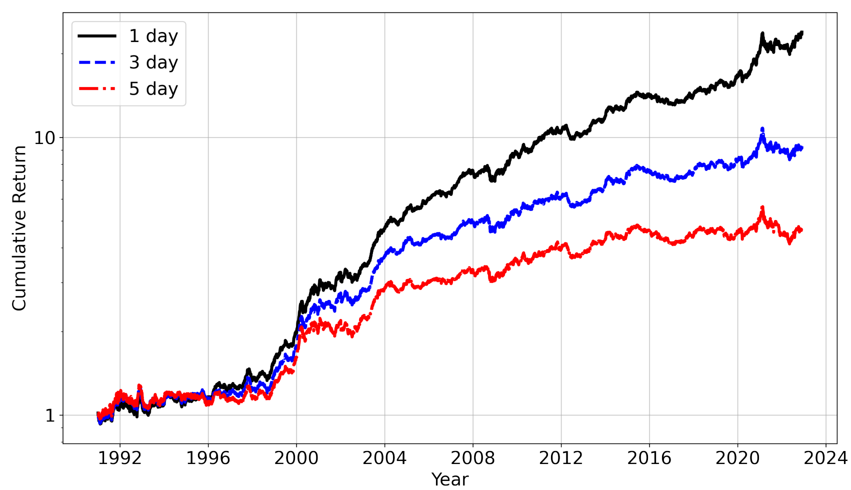

Figure 6: Conditional Probability an Insider's Trade is Informed

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}, whether the trade was a purchase or sale, and attributes of the insider's past trading history. Specifically, the probability is conditioned on the insider's past average abnormal return $\bar{r}_i$ and the standard deviation of their previous returns σ_i. The number of past trades is held fixed at ten past trades. Each panel shows conditional probability curves for the standard deviation in the panel header and for past average returns of -2%, 0%, and 2%. The top (bottom) row shows the curves for purchases (sales).

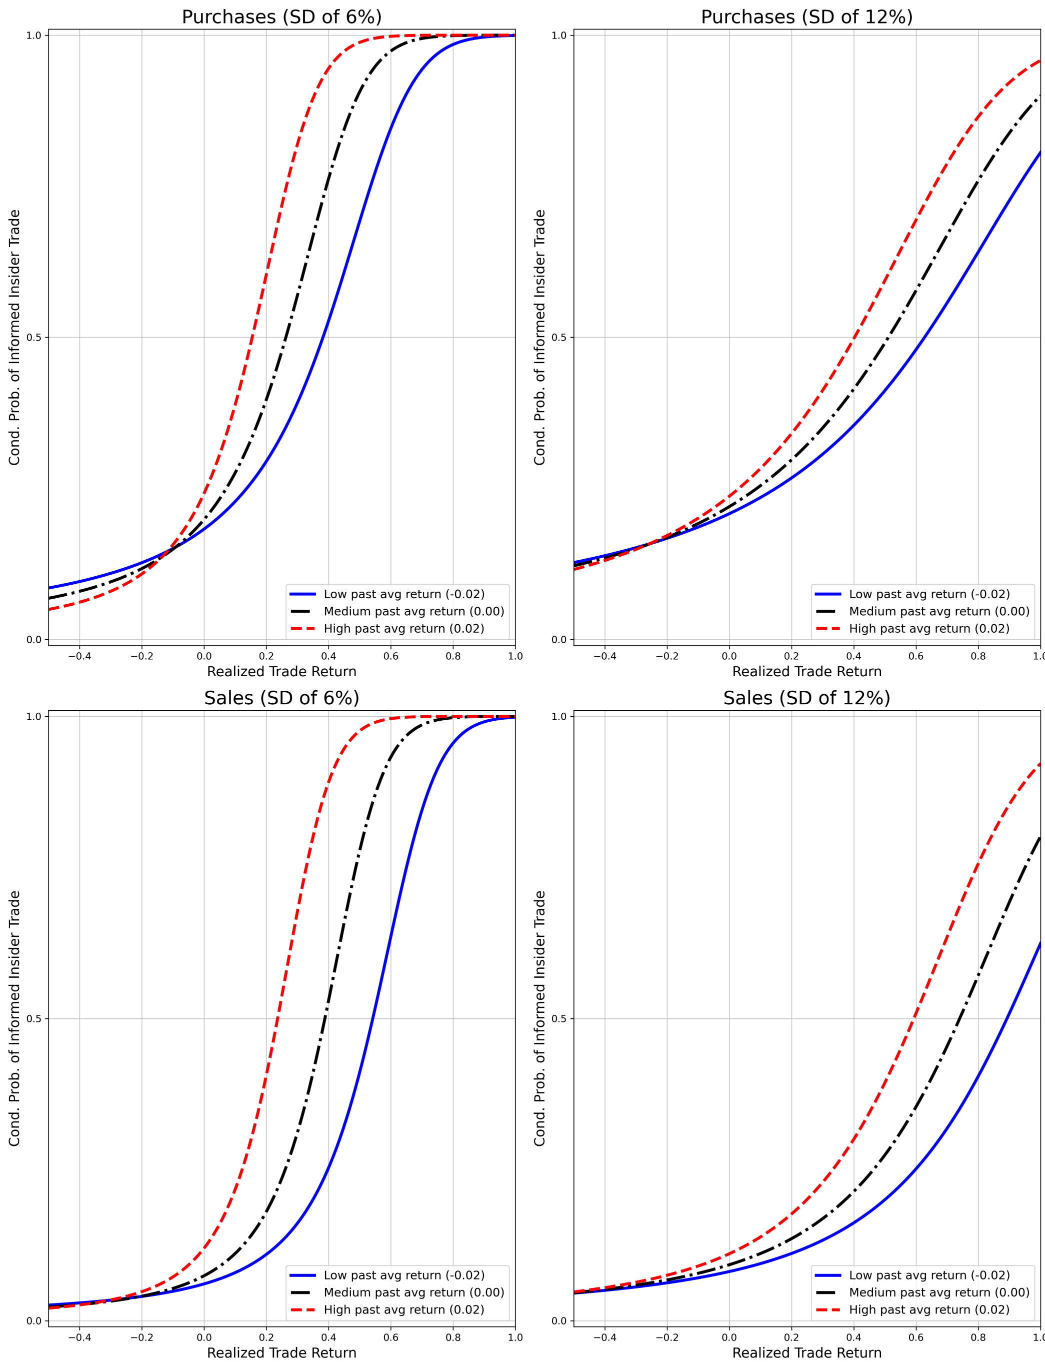

Figure 7: Insider-specific Return Thresholds

The figure plots the trade-level signed abnormal return threshold at which it is more likely than not that a trade made by a corporate insider was informed. The threshold depends on the insider's past average abnormal return $\bar{r}_i$, whether the trade was a purchase or sale, the standard deviation of their previous returns σ_i, and the number of past trades. The top (bottom) row shows the curves for purchases (sales). The number of past trades is held fixed at ten past trades. The thresholds are plotted as a function of the insider's past average return for three levels of past standard deviation of past abnormal trade returns. The low, median, and high noise levels correspond to the 10th, 50th, and 90th percentiles of trade-level standard deviations for a given (binned) past average return.

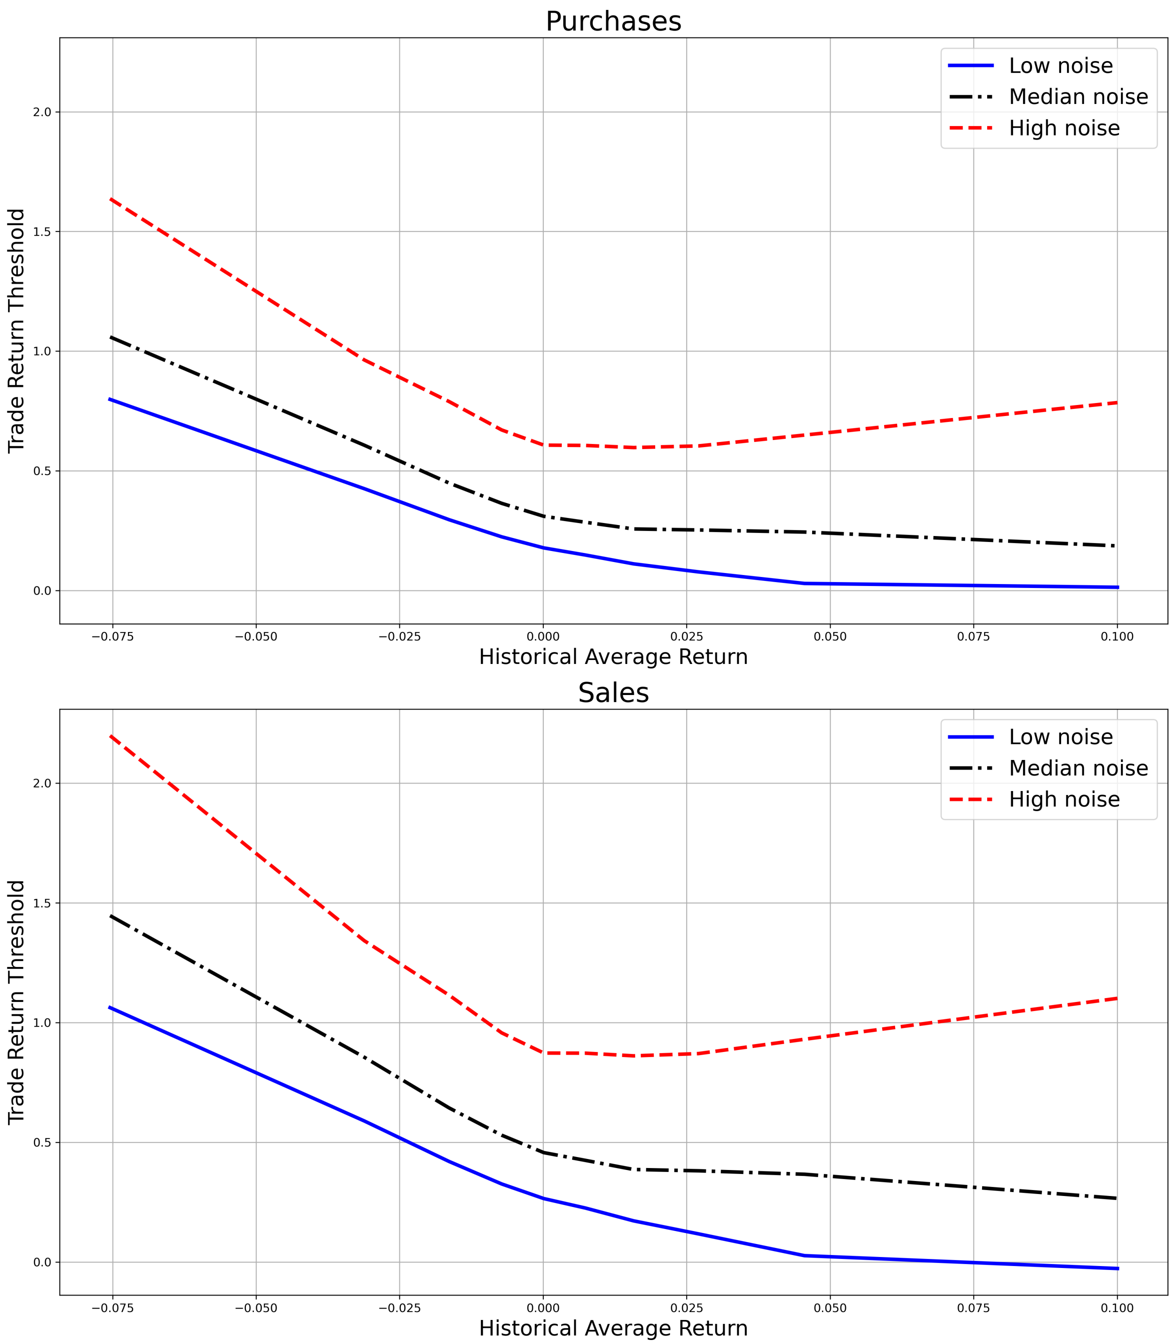

Figure 8: Case Study: Enron

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

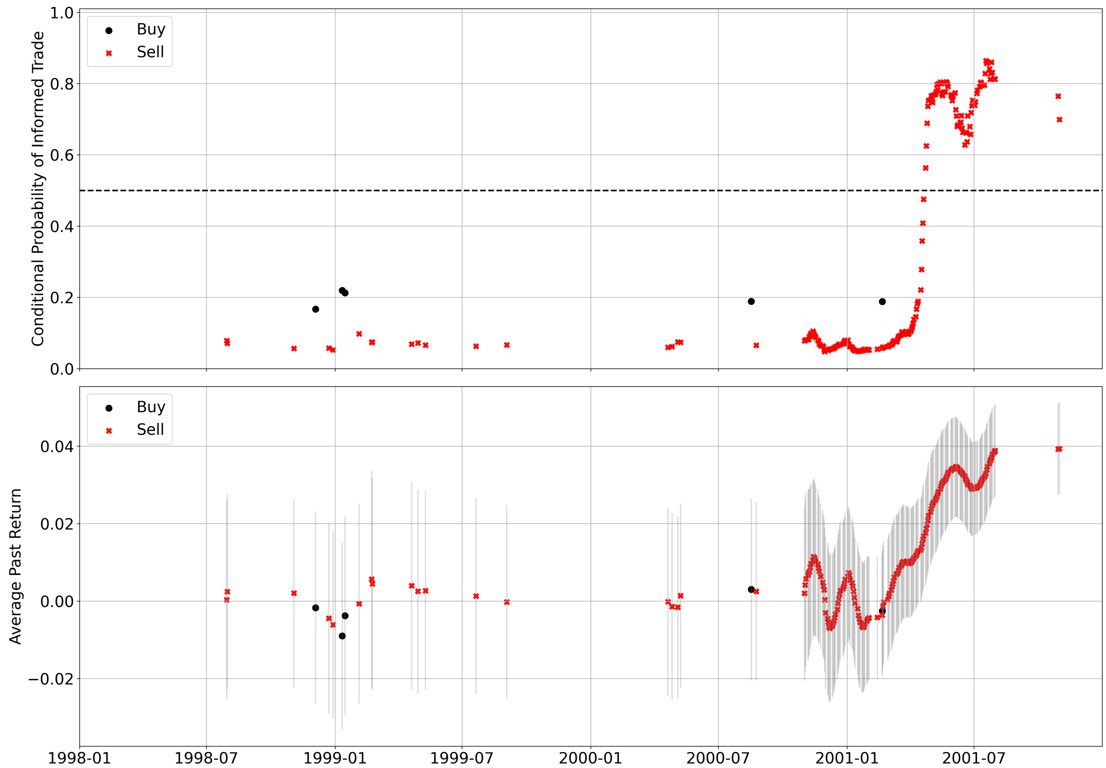

(b) Jeffrey Skilling