Conviction and volume: Measuring the information content of hedge fund trading

Jonathan Rhinesmith*

December 30, 2015

Abstract

I provide novel evidence that hedge funds predict and drive the movement of asset prices towards fundamental value. Willingness to move prices, proxied by the share of trading volume consumed, reveals information: the volume consumed by quarterly hedge fund trades strongly predicts future stock returns. The top decile of purchases generates abnormal returns of 5-9% annualized during the following quarter (t-stat 4.4-6.5). Interpreting this phenomenon using the Kyle model of price impact, I test for the empirical patterns one should observe if informed (hedge fund) trades incorporate information into prices. Informed trading impounds earnings news, reducing the reaction to positive earnings announcements by 28%. Informed trading also positively predicts contemporaneous price movement and future informed trading. These price movements do not reverse. In contrast, mutual fund trades are significantly less informative. Structural and reduced-form estimates imply that consuming 1% of quarterly volume generates 0.3%-0.5% of price impact. Taken together, these results suggest that funds incorporate substantially more information into prices than is apparent from their fund-level returns.

*Harvard University, rhinesmith@fas.harvard.edu. I am particularly grateful to John Campbell, Jeremy Stein, Lauren Cohen, and Chris Malloy. I also thank Malcolm Baker, Robin Greenwood, Sam Hanson, Owen Lamont, Jonathan Reuter, Andrei Shleifer, Erik Stafford, Adi Sunderam, and seminar participants at Harvard and the Federal Reserve Board of Governors for helpful discussions and suggestions. I thank Vikas Agarwal, Vyacheslav Fos, and Wei Jiang for providing me with data. For the latest version, please visit: scholar.harvard.edu/rhinesmith/jobmarketpaper

Introduction

In this paper I study hedge fund trading with two questions in mind. First, are hedge funds informed? Second, if so, how does their information get incorporated into prices? I show that trading volume plays a key role in addressing these questions.

I apply the intuition of microstructure models – which are typically considered at daily horizons – to the quarterly investment behavior of hedge funds. This approach provides novel insights into the above questions. In particular, I draw on the intuition of the Kyle (1985) model that price impact is a function of volume. An informed fund should trade until the marginal cost of price impact equals the marginal profit of trading an additional share. Willingness to move prices reveals information: if large trades relative to volume cause price impact, then fund managers should only be willing to consume a large share of volume when their private information is especially compelling. Following this logic, I study the "volume consumed" – shares traded divided by total volume – by quarterly hedge fund trades.

I demonstrate that the cross section of volume consumed strongly predicts stock returns during the following quarter. The top decile of hedge fund equity purchases by volume consumed generates statistically significant outperformance of 5-9% annualized during the following quarter (t-stat 4.4-6.5). The top five deciles of purchases, representing 79% of purchases by dollar value, display statistically significant outperformance. I focus on purchases because I observe hedge funds' long portfolios.[1] These results suggest that hedge funds are informed.

To study how this information gets into prices, I test for the empirical patterns one should observe if the price impact of hedge fund trades incorporates information. Informed trades prior to the public revelation of earnings should impound earnings information into prices. The associated stocks should then react less when earnings news is revealed. Confirming this reasoning, I find that the reaction to a given positive standardized unexpected earnings surprise (SUE) is reduced by 28% for stocks in the top quintile of volume consumed relative to stocks with no hedge fund activity. I study positive surprises because of my focus on the information content of purchases. Though hedge fund purchases reduce the returns associated with a given earnings surprise, purchases nevertheless predict earnings returns

[1]For both empirical and theoretical reasons, within a long portfolio, purchases are more likely to convey private information than sales. Chan and Lakonishok (1993): "Information effects might also be stronger for purchases than for sales...[The] choice of a particular issue to sell, out of the limited alternatives in a portfolio, does not necessarily convey negative information. Rather, the stocks that are sold may already have met the portfolio's objectives, or there may be other mechanical rules, unrelated to expectations about future performance, for reducing a position...In contrast, the choice of one specific issue to buy, out of the numerous possibilities on the market, is likely to convey favorable firm-specific news."

unconditionally (before controlling for the level of the earnings surprise).

I provide three more important pieces of evidence that hedge fund trades incorporate information into prices. First, I show that the prices of high volume-consumed positions increase as hedge funds buy them. This pattern is consistent with price impact.

Second, I show that trading is persistent across time. Purchases in quarter t predict purchases in quarter t+1. In quarter t+1, funds buy a greater share of volume in stocks with high quarter t volume consumed than in stocks with low quarter t volume consumed. During quarter t+1, the former positions perform better than do the latter positions. If funds do not cause price impact, then they are leaving money on the table by not building the former positions even faster.

Third, the cumulative outperformance of high volume-consumed positions is significantly positive out to a horizon of 2-4 years. Hedge fund trading is associated with fundamental information, which I define as persistent long-horizon price movements, rather than temporary price pressure, which would revert. This test rules out the possibility that hedge funds merely predict the price impact of their own future trades.

These results are based on trades identified from 13F filings. My hedge fund sample captures \$200 billion of equity positions at a given time, on average, and over \$500 billion by the end of the sample. The data covers \$4.3 trillion of purchases, 1.0% of total volume.

In contrast to large hedge fund trades, mutual fund trades that are large relative to volume are significantly less informative. Large mutual fund trades generate strong contemporaneous performance. Trades should cause price impact as they occur, regardless of information content. However, these trades predict at best marginally positive future performance, even after removing funds subject to extreme fund-level flows. This performance tends to revert over long time horizons, which should only occur for non-information-based trades.[2]

Yet there is evidence that a subset of mutual funds are skilled. If informed volume reveals information, then volume consumed within this subset should predict future returns. I confirm this prediction using measures of skill from the literature.[3]

I derive these tests from a two-period Kyle model, which intuitively formalizes how informed trades impound information into prices. I treat calendar quarters as periods. An informed trader balances her price impact – which reduces profits per share – against a desire to trade more shares – which increases the quantity she profits on. She also balances how much to trade this quarter against how much to trade next quarter. Trading over multiple quarters reduces the effect of price impact but increases the risk that information will be

[2]Coval and Stafford (2007), Frazzini and Lamont (2008), and Lou (2012) find reversals following mutual fund flow-driven trades.

[3]Specifically, I examine return gap (Kacperczyk, Sialm, and Zheng (2008)) and active share (Cremers and Petajisto (2009)).

publicly released before the trader has finished building her position.

The model also generates quantitative, parametric implications for the comovement of trading and prices. The optimal amount of informed trading is linearly related to a stock's mispricing times expected noise trading. Furthermore, permanent price impact is linear in trade size. The price impact function is of interest because it captures how much information a given amount of informed trading incorporates into prices. I structurally estimate the model using maximum likelihood and also directly estimate the model-implied reduced form for price impact using Fama-MacBeth regressions. I find that purchasing 1% of the volume in a stock over a calendar quarter generates 0.3%-0.5% of permanent price impact. The structural model generates simulated moments of trading and returns that are reasonably close to the corresponding empirical moments.

My findings connect to the literatures on informed trading, active investment management, and market microstructure. Hedge funds are not the only market participants with differential information. For instance, firm insiders may be particularly well informed about a company's prospects, and insiders' trades are known to predict equity returns. I show that hedge fund volume consumed comoves positively with the purchases of firm insiders. However, insider trading does not subsume hedge fund purchases.

On the other hand, the literature suggests that investment funds may reveal information through channels other than volume. In a world without trading costs, a fund manager should trade on a piece of private information until she hits a risk limit: a fund's largest risk-weighted positions should have the highest expected returns. In my sample, however, I find that volume consumed subsumes idiosyncratic risk-weights.

My results suggest that if trades generate price impact, then one must examine asset prices before a fund's first trade to properly account for the information that a fund incorporates into prices. Funds move prices as they build large positions. Neither the post-purchase prices of investment holdings nor fund-level returns – two metrics that the literature often focuses on for other purposes – fully account for this effect. A fund with poor returns based on these metrics could still be identifying a substantial amount of information and helping to incorporate that information into prices.

Long-horizon price impact also contributes to decreasing returns to scale in active management, as in Berk and Green (2004) and Pastor and Stambaugh (2012). My findings offer a quantification of how quarterly trades generate price impact at the individual stock level. Existing evidence for diseconomies tends to focus on the fund and industry levels.

Finally, the market microstructure literature provides evidence of price impact at intraday and daily horizons. I present evidence of price impact at quarterly time horizons. Quarterly price movements are more relevant to many of the economic decisions of firm managers. The disadvantage of moving to a coarser time horizon is that causation is not as clear-cut. I rely on the plausible assumption that hedge funds scale trades optimally given information – that they do not systematically leave money on the table – to rule out the possibility that price movements are exogenous. Without price impact, the tendency of funds to size trades relative to volume and to trade in a persistent manner would be suboptimal.

The paper proceeds as follows. Section 1 reviews related literature. Section 2 develops a Kyle model of price impact and uses the model to generate testable hypotheses. Section 3 describes the data and constructs volume consumed. Section 4 presents the core empirical results of the paper, evidence that hedge fund volume consumed reveals information and incorporates some of that information into prices. Section 5 shows that total mutual fund volume consumed is uninformative but that plausibly skilled subsets of funds reveal information through volume. Section 6 estimates the quantitative price impact function, both by employing the reduced form and by undertaking a structural estimation of the model. Section 7 shows that hedge funds trade alongside firm insiders and that volume consumed subsumes idiosyncratic risk-weights in my sample. It also considers a publicly implementable trading strategy. Section 8 concludes.

1 Literature

An extensive literature examines skill in the active management industry. Superior net-of-fee mutual fund returns are difficult to consistently identify (Fama and French (2010)). Del Guercio and Reuter (2014) show that broker-sold retail mutual funds exhibit negative post-fee returns, while direct-sold retail funds exhibit post-fee returns indistinguishable from zero. Hedge fund net-of-fee skill is also subject to debate. Properly adjusting for the risk of funds' returns is made more complicated by the use of options (e.g., Jurek and Stafford (2015)) and possible reporting biases (e.g., Patton, Ramadorai, and Streatfield (2015)). Studying long U.S. equity holdings allows me to employ standard risk-adjusted equity returns.

There is some evidence of gross-of-fee skill in subsets of hedge funds and mutual funds (or subsets of their holdings). Berk and van Binsbergen (2015) find evidence of mutual fund skill using gross dollar value added. Griffin and Xu (2009) and Agarwal, Fos, and Jiang (2013) find that hedge funds demonstrate weakly positive gross skill on their overall equity holdings. Cohen, Polk, and Silli (2010) show that mutual funds outperform on their largest risk-weighted positions. Cremers and Petajisto (2009) examine mutual funds' deviations from benchmark weights. Cohen, Frazzini, and Malloy (2008) find outperformance in stocks where mutual fund managers share an educational connection with board members. Rhinesmith (2014) documents that hedge funds outperform in the stocks that they "double down" on after poor stock-level performance.

Other mutual fund trades appear to drive price dislocations and subsequent long-horizon reversals. Coval and Stafford (2007) and Frazzini and Lamont (2008) provide evidence based on fund-flow-driven trades. Lou (2012) links fund flows to momentum. Khan, Kogan, and Serafeim (2012) and Dasgupta, Prat, and Verardo (2011) find reversals following general large mutual fund purchases and institutional herding, respectively. My hedge fund findings contrast with these papers, as I show long-horizon outperformance.

Hong, Li, Ni, Scheinkman, and Yan (2015) show that short ratio divided by volume predicts future (negative) stock returns better than the unadjusted short ratio. In contrast to this paper, Hong et. al. focus on the short side. It is difficult to break down short ratios across different investors. The authors focus less on how information gets into prices, or on long-horizon returns. Finally, their measure takes the level of short interest and divides it by volume, since holding a short position faces a stock-lending friction. I focus on trades (changes). There is no clear friction that inhibits holding a long position.

An emerging literature examines how hedge funds impact equilibrium prices. Kruttli, Patton, and Ramadorai (2014) show that hedge fund illiquidity forecasts the returns of equity, bond, and currency portfolios. Cao, Chen, Goetzmann, and Liang (2015) make the case that hedge funds hold stocks that are above the security market line, and that hedge fund ownership precedes the dissipation of alphas. I focus more on the mechanism through which hedge funds eliminate mispricings.

The market microstructure literature documents that trading appears to incorporate private information into prices. Kyle (1985) lays out a workhorse model of how trading volume and prices are determined in equilibrium. I review this model in Section 2. Holden and Subrahmanyam (1992), Foster and Viswanathan (1996), and Back, Cao, and Willard (2000) analyze the model with multiple informed traders. Huberman and Stanzl (2004) generalize the linear relationship of permanent price impact and trade size.

Koudijs (2014) provides empirical support for the Kyle model in a natural experiment with identifiable private information. He examines the comovement of the returns of dual-listed stocks and news arrivals. Boulatov, Hendershott, and Livdan (2013) study cross-asset implications of the model. These studies focus on time horizons of a handful of days.

In contemporaneous research, Di Mascio, Lines, and Naik (2015) analyze a proprietary dataset and also find that institutions trade in the same direction in the same stock over multiple quarters, and that only purchases are informative (not sales). Their different frequency (daily) and sample (a mix of institutions and equity markets) provide complementary evidence for these two findings. I focus more on price impact as a function of volume. My paper considers how skill differs by investor type and examines longer horizon returns. Much of the evidence for reversals following institutional trades occurs at a multi-year horizon.

Another literature focuses on empirical transaction costs. In seminal papers, Keim and

Madhavan (1995, 1996) and Chan and Lakonishok (1993) examine the intraday price impact of institutional trades. In a similar spirit to my work, Chan and Lakonishok (1995) study the combined price impact of packaged trades, but at much shorter time horizons.

Other studies examine further the relationship of institutional trades to returns. See Campbell, Ramadorai, and Schwartz (2009) for an overview.[4] Several papers find that quarterly institutional flows are positively correlated with contemporaneous stock returns. These papers do not separate out hedge fund trading, which may be differentially informative. This paper also focuses more on trading volume (relative to quarterly horizon studies) and finds stronger evidence of long-horizon return persistence than much of the literature.

Industry anecdotes confirm that price impact considerations could lead a fund to stagger trading in a single stock over more than one quarter. In the literature, hedge funds attach a high value to delaying their 13F filings. Agarwal, Jiang, Tang, and Yang (2013) show that confidential 13F holdings strongly outperform. This result suggests that funds worry that other market participants may try to frontrun them at a quarterly frequency.

Berk and Green (2004) and Pastor and Stambaugh (2012) present seminal models of the active management industry. Both papers assume decreasing returns to scale – at the fund and industry level, respectively – as a linchpin of their models. Chen, Hong, Huang, and Kubik (2004), Pollet and Wilson (2008), and Pastor, Stambaugh, and Taylor (2015a, 2015b) empirically examine returns to scale at the mutual fund and industry levels. This paper provides evidence at the stock / trade level for diminishing returns to scale. Existing evidence at this level of disaggregation is typically based on short time horizons (such as the microstructure literature) or special cases (such as firesales and extreme fund flows).

2 Two-period Kyle model

The Kyle model intuitively formalizes how the price impact of informed trades impounds information into prices. I construct a two-period version of the model to generate my hypotheses. I treat quarters as periods. Relative to other models of informed trading and investing, the Kyle model makes three key points. First, it focuses on volume, particularly informed volume relative to uninformed volume. Second, it considers specifically permanent price impact, which reflects information. Third, it suggests that private information is best recovered from changes rather than from levels, or from informed trades rather than from informed holdings. All proofs are detailed in Appendix A.

[4] Notable additions since then include Choi and Sias (2012), Hendershott, Livdan, and Schurhoff (2015), and Collin-Dufresne and Fos (2015).

In the model, an informed trader possesses a piece of private information that will be publicly released after either the first or second quarter. The informed trader chooses how many shares of stock to trade each quarter in order to maximize profits. Concurrently, a random amount of noise trading arrives each quarter. A competitive market maker observes total net order flow (informed plus noise trades), and conditional on that observation sets the stock's price equal to its expected value. The market maker absorbs the net order flow at that price. The market maker and informed trader are risk neutral. In expectation, the informed trader profits, the market maker breaks even, and the noise trader takes losses.

The informed trader balances her price impact (less profit per share) against a desire to trade more shares (a higher quantity on which to profit). She also balances how much to trade this quarter against how much to trade next quarter. Trading over a longer period of time reduces the average price impact of trades, because new noise traders arrive each period. However, trading slower increases the risk that the information will be released before the informed trader finishes building her desired position.

I model a single informed trader to capture the key intuition that drives my hypotheses. While I observe multiple hedge funds, how the model's implications change as one varies the number of informed traders depends on further assumptions. I briefly discuss this point following my hypotheses.

Formally, assume that a single risky asset receives a piece of new information at the start of a two-quarter information "episode." There are four deep parameters: the variances of (1) information (σ_ϵ^2), (2) the noise in the insider's signal (σ_η^2), and (3) noise trading (σ_u^2), as well as (4) the probability that information will be released late (π). Assume the asset has a price of 0 at the beginning of the current episode. With the new information, the asset is worth $\epsilon \sim N(0,\, \sigma_\epsilon^2)$. The informed risk neutral insider observes the true information plus noise, or $i = \epsilon + \eta$, where $\eta \sim N(0,\, \sigma_\eta^2)$ and is independent of all other random variables. The expectation of information given her signal is ϕi, with $\phi = \frac{\sigma_\epsilon^2}{\sigma_\epsilon^2 + \sigma_\eta^2}$. To all market participants, ϵ is publicly revealed shortly after the end of quarter 1 (early) with probability $1 - \pi$ and shortly after the end of quarter 2 (late) with probability π.

The insider trades an amount x_t each period. Notably, x_2 occurs only if the information is not revealed early. The Kyle model does not describe trading after (or during) the information revelation event. Noise traders also arrive each period, with noise trading $u_t \sim N(0,\, \sigma_u^2)$, independent of ϵ.[5]

The competitive risk neutral market maker observes total order flow $y_t = x_t + u_t$. The market maker can go long or short, and sets price p_t at the conditional expected value of ϵ.

[5]It is straightforward to allow the variance of noise trading to differ across the early and late periods. The main conclusions of the model carry through.

That is, $p_t = E[\epsilon \,|\, Y_t]$, where $Y_1 = \{y_1\}$ and $Y_2 = \{y_1,\, y_2\}$.

Once the information is revealed, the price immediately equals fundamental value. Permanent price impact is defined with respect to information: it is price impact that persists until the public revelation of information. In the sense of the model, it is price impact that reflects the (average) private information revealed through trading. Temporary price impact reverts following a large trade, even in the absence of the release of information.

The model has a unique linear equilibrium. The insider trades an amount each period that is linearly related to the remaining mispricing:

$$x_1 = \beta_1 \phi i \tag{1}$$
$$x_2 = \beta_2(\phi i - p_1) \tag{2}$$

The market maker in turn sets prices as a linear function of total order flow:

$$p_1 = \lambda_1(x_1 + u_1) \tag{3}$$
$$p_2 = p_1 + \lambda_2(x_2 + u_2) \tag{4}$$

The equations determining the parameters λ_1, λ_2, β_1, and β_2 are detailed in Appendix A.1. Those equations are developed as equilibrium constraints from the informed trader's maximization problem and from the market maker's inference of expected value conditional on total order flow.

My hypotheses follow in chronological order. For reference, Figure 1 displays these hypotheses on a timeline.

Hypothesis (1) *More informed trading in quarter 1 implies more price movement in quarter 1 [$cov(p_1,\, x_1) > 0$].*

Trades cause price impact. If the insider trades more this quarter, the price will move by more this quarter.

Hypothesis (2) *More informed trading in quarter 1 implies more price movement in quarter 2 [$cov(p_2 - p_1,\, x_1) > 0$ and $cov(\epsilon - p_1,\, x_1) > 0$].*

If there is a larger initial mispricing, the insider trades more this quarter (x_1 is linearly related to ϵ plus noise). In addition, the price will move by more next quarter. If the information is not revealed early, the insider will push the price by trading more next quarter (relative to reduced-information trades). If the information is revealed early, then the revelation of information will move price by a greater amount (because the initial mispricing was greater).

Hypothesis (3) *More informed trading in quarter 1 implies more informed trading in quarter 2, if the information is not revealed early [*$cov(x_2, x_1) > 0$*].*

If information is not revealed early, then the informed insider will continue to trade more next quarter than she would for a smaller initial mispricing.

Hypothesis (4) *Given fixed true positive information (*$\epsilon > 0$*), more informed trading in quarter 1 implies a smaller price reaction when the information is revealed [*$\frac{\partial}{\partial \eta'} E(\epsilon - p_1 | \eta = \eta', \epsilon > 0) < 0$*].*

For a positive true information draw $\epsilon > 0$, suppose the insider buys more in quarter 1 as a result of the noise in her signal, η. In that case, she will incorporate more of the information into the price prior to the information's public release. When that information is then revealed, the price will react less. (The model has a symmetric implication for negative news.)

A new information event is assumed to occur every two quarters. The risky asset begins each episode with its price equal to its fundamental value at the end of the prior information episode (rather than zero). For clarity, I only utilize the formal notation required to track prices over time when referencing the following hypothesis. Denote the fundamental value of the asset after K episodes as $\bar{\epsilon}^K = \sum_{k=1}^{K} \epsilon_k$, with ϵ_k equal to the information draw from episode k, and assume that we take the perspective of the first information event.

Hypothesis (5) *Price movement over the course of the current information episode persists into the future [*$cov(\bar{\epsilon}^K, x_1) > 0$*,* $K > 1$*].*

The future price path subsequent to the current information episode follows a martingale (future information draws are mean zero). Cumulative performance over the course of the current episode should persist (no reversals). In this sense, prices move to reflect information.

As presented above, total informed volume emerges as a summary statistic for private information in the version of the model with a single informed trader. In a model with a varying number of informed traders, it is no longer generally the case that total informed volume is a summary statistic. However, the generalization depends on the details. For example, take Hypothesis (2), which states that informed volume in quarter 1 predicts returns in quarter 2. With perfectly correlated signals, an increase in the number of informed agents for a fixed total amount of informed volume reflects increased competition rather than increased information, reducing expected returns in quarter 2. In contrast, if signals are imperfectly correlated and if the econometrician only observes purchases (as in my data), then observing multiple informed agents validates the information and can increase the expected value of the asset. In forming the posterior of the information, two observations shift the prior further than one does. This effect can increase expected returns in quarter 2. I elaborate on this point and discuss competition in more detail in Appendix F.

3 Constructing volume consumed

3.1 Data

I construct my sample by linking the Thompson Reuters database of publicly available Form 13Fs, which contain the quarterly holdings of asset management institutions, to a sample of hedge funds identified by Agarwal, Fos, and Jiang (2013).

I begin with the Thompson Reuters 13F database. Any investment management institution that "exercises investment discretion over $100 million or more in Section 13F securities" (generally long U.S. equity positions, as well as some derivatives) is required to file a 13F within 45 days of the end of every calendar quarter.[6] The Form 13F reports the list of 13F securities that the investment manager holds as of the end of the corresponding quarter.[7] I focus on the sample of the 92 13Fs filed between 12/31/1989 and 9/30/2012.

I identify hedge funds using the comprehensive set of funds from Agarwal, Fos, and Jiang (2013). As explained in more detail in their paper, the authors merge five large commercial hedge fund databases with industry publications to form their dataset.

I obtain stock return and volume data from CRSP, and stock accounting data from Compustat. I focus on common stocks (CRSP share codes 10 and 11). I use the procedure of Shumway (1997) to account for delisting returns. I construct characteristic-adjusted returns following the procedure of Daniel, Grinblatt, Titman, and Wermers (1997).[8] The characteristic-adjusted return of a stock is the return of that stock minus the value-weighted return of a portfolio of stocks matched to have the same size, value, and momentum characteristics as the stock in question (using a 5x5x5 sort to produce 125 matching portfolios). Market-adjusted returns subtract the returns of the CRSP value-weighted index. Risk factor returns (SMB, HML, UMD) are from Ken French's website. I use insider trading data from the Thompson Reuters database of Form 4s. Analyst estimates are from I/B/E/S. Mutual fund data is from the Thompson Reuters mutual fund holdings database, as well as the

[6] More detailed requirements are provided at https://www.sec.gov/answers/form13f.htm. The full list of 13F securities is available at http://www.sec.gov/divisions/investment/13flists.htm.

[7] As is common in the literature, I ignore the 45-day filing delay. Instead, I analyze holdings as of the date the manager holds those positions. This approach focuses on the behavior of fund managers, rather than attempting to construct a trading strategy that can be implemented by a third-party using publicly available information. In Section 7.3, I conduct a brief analysis that incorporates the filing delay.

[8] The DGTW benchmarks are available via http://alex2.umd.edu/wermers/ftpsite/Dgtw/coverpage.htm

CRSP survivorship free mutual fund database.

I seek to infer information from the behavior of hedge funds. Unfortunately, 13F filings do not provide information on short positions, cash holdings, or non-U.S. equity positions. I therefore remove filings that are unrepresentative of a firm's active investment strategy. For example, a fund that reports only a single stock on a Form 13F is probably investing primarily outside of publicly listed U.S. equities. Other fund companies hold a disproportionately large number of stocks. It is relatively more likely that these 13F filings encompass multiple underlying funds, potentially based on different investment strategies. Overlapping funds make it more difficult to infer information-based trades. Each underlying fund can be subject to asset flows of different relative magnitudes, and trades can cross. It is also difficult to estimate the percent of a given fund's portfolio that any position represents, which I use in some tests. A large position in a small underlying fund could show up as a small percent of a firm's total 13F portfolio. Furthermore, the 13F filings of very diversified hedge fund firms are more likely to reflect index-relative investment strategies. Some firms allow individual clients to customize their benchmark indices. This approach makes it difficult to separate index-tracking trades from active trades based on private information, since a corresponding 13F filing aggregates multiple client portfolios that track separate indices.[9]

I therefore remove (1) any filing with fewer than 10 positions, (2) any filing which contains more than 150 positions, and (3) any filing in which the value of the 13F portfolio is under \$50 million. The mutual fund literature employs similar standard screens.[10] When industry practitioners analyze the information content of 13F portfolios, they also eliminate filings based on a minimum and maximum number of 13F positions. The implication is that practitioners believe the remaining 13Fs are the most informative.[11] None of my results are sensitive to these particular threshold values. These screens reduce my sample of fund-quarters from 44,126 observations to 28,128 observations.[12] All discussion of manager returns, trades, flows, and relative position sizes refer to the remaining 13F portfolios.

Table 1 panel A summarizes the hedge fund universe across the 92 13F filings in my sample. Averages are taken in the time series. Overall, hedge funds hold large cap stocks

[9]For example, D.E. Shaw, one of the largest quantitative hedge fund firms, "enables...investors to customize their exposure to a particular index" in some funds (from http://www.deshaw.com/WhatWeDo.shtml, accessed 8/10/2015).

[10]See, for instance, Kacperczyk, Sialm, and Zheng (2008).

[11]Goldman Sachs' "Hedge Fund Trend Monitor," for example, removes funds with fewer than 10 positions or more than 200 positions. It describes this requirement as "an attempt to isolate fundamentally driven investors from quantitative funds or funds that mirror private equity investments" (from http://www.bloomberg.com/news/articles/2015-05-21/goldman-these-are-the-100-most-important-stocks-to-hedge-funds, accessed 5/22/2015).

[12]Two examples of funds removed by this procedure at 9/30/2012 are TA Associates and D.E. Shaw. TA Associates is a private equity firm, which listed only two positions on its 13F for 9/30/2012. D.E. Shaw is a quantitative hedge fund firm that held 1,783 positions at 9/30/2012.

with slight value and momentum tilts. The sample grows steadily over time, and peaks at 572 managers in late 2007. This hedge fund universe captures \$200 billion of long equity positions, on average, or roughly \$500 billion by the end of the sample. My data covers \$4.3 trillion of purchases in total, 1.0% of overall equity market trading over the sample period.

3.2 Construction of volume consumed

In the Kyle model, the optimal amount of informed trading is linearly related to a stock's mispricing times expected noise trading. In a one-period Kyle model, this result is trivial. I show in Appendix A.8 that this result also holds in my two-period model: $x_1 = constant * E[mispricing] * \sigma_u$, where x_1 is informed trading and σ_u measures the magnitude of expected noise trading (the expectation of the absolute value of a mean-zero normal random variable is proportional to its standard deviation). A similar result holds for x_2.

If one believes that hedge funds may be informed, then this optimum suggests an observable proxy for mispricing based on volume: hedge fund (informed) volume divided by a stock's normal (uninformed) volume.[13] To implement this proxy, I construct a measure of the volume consumed by a given hedge fund in each of its individual stock positions, relative to lagged volume (in shares): $volconsumed^{\pm}_{s,f,t} = \frac{shares\,traded_{s,f,t}}{volume_{s,t-1}}$, in stock s, for fund f, during quarter t, where $shares\,traded_{s,f,t}$ is positive for purchases and negative for sales.[14] In order to deal with trends and seasonality in volume, when constructing portfolios I ensure that all positions being sorted use volume and trades measured over the same time periods.

Volume consumed: $volconsumed^{\pm}_{s,f,t} = \frac{shares\,traded_{s,f,t}}{volume_{s,t-1}}$,

$$volconsumed_{s,f,t} = max(volconsumed^{\pm}_{s,f,t}, 0)$$

I primarily focus on purchases, or $volconsumed_{s,f,t} = max(volconsumed^{\pm}_{s,f,t}, 0)$. Purchases in a long portfolio are more likely to reflect information than sales of existing positions. A fund manager chooses to purchase a stock from thousands of listed stocks, whereas when she needs to sell a stock to generate cash for outflows or new investments, or to reduce market exposure, she chooses from her limited set of existing holdings. I cannot see hedge fund short positions, which censors negative observations of volume consumed. Short sale constraints also suggest that the relationship between volume consumed and information may differ when sales cross a zero position level. Inferring information from negative volume

[13]Treating "normal" volume as a proxy for noise trading follows in the tradition of microstructure models. Easley, Hvidkjaer, and O'Hara (2002) note that their model "in effect...interprets the normal level of buys and sells in a stock as uninformed trade."

[14]I use lagged volume because a fund may not know contemporaneous volume before it decides to trade. My core results are robust to using contemporaneous volume, contemporaneous volume excluding hedge fund trades, or more distant lags of past volume.

consumed is difficult, even if some sales are information-driven.

From another perspective, consider a fund that possesses perfect private information and maximizes expected returns. In a frictionless world without bubbles it is not clear why such a fund would *ever* hold an overvalued position in its long portfolio for a non-zero amount of time. It would buy undervalued stocks and sell them once they reached fair value. In such a scenario, sales of long holdings would not predict future underperformance.

I construct volume consumed using three different methods. I primarily focus on the first construction. This construction – aggregation method 1 – sums all purchases at the stock level to produce an aggregate amount of hedge fund purchase volume. It does not net out sales. That is, $volconsumed_{s,t} = \sum_{f=1}^{F} volconsumed_{s,f,t}$ for a stock s and a quarter t, summing $volconsumed_{s,f,t}$ in stock s across all funds F. I primarily focus on this method because the price impact of multiple purchases in a single stock should aggregate.

In the second construction, I look at purchases at the manager-stock level. This method produces more observations than method 1, since a single stock could be purchased by multiple managers in a given quarter. In aggregation method 2, I have $volconsumed_{s,f,t}$ triplets for a stock s, a fund f, and a quarter t.

The third construction nets purchases and sales. Aggregation method 3 produces a single net amount bought (positive) or sold (negative) by all hedge funds F in a stock s during quarter t: $volconsumed^{\pm}_{s,t} = \sum_{f=1}^{F} volconsumed^{\pm}_{s,f,t}$. Sales are uninformative, so I primarily utilize this approach for robustness.

Aggregation method 1 (my focus): $volconsumed_{s,t} = \sum_{f=1}^{F} volconsumed_{s,f,t}$
Aggregation method 2: $volconsumed_{s,f,t}$
Aggregation method 3: $volconsumed^{\pm}_{s,t} = \sum_{f=1}^{F} volconsumed^{\pm}_{s,f,t}$

To form volume consumed quintile or decile portfolios, I sort separately among (1) NYSE/Amex and (2) NASDAQ stocks, and then combine the resulting portfolios. Historically, NASDAQ volume figures are not comparable to NYSE/Amex volumes.[15]

I present most of my results using Fama-MacBeth regressions. A regression summarizes the relationship of volume consumed and a given variable using a single coefficient. For my tests that focus on investment performance, however, I present results that are based on forming explicit calendar-time portfolios. Portfolios facilitate comparison with the literature (future returns, Hypothesis (2)) and ensure proper standard errors (long-horizon returns, Hypothesis (5)). Portfolio results for other tests are available in Appendix E.

Hedge fund information is not limited to small cap stocks with low volume. To emphasize this fact, in my portfolio results that use volume consumed in quarter t to predict returns

[15]For more details, see Anderson and Dyl (2005).

in quarter t+1 – a standard test of skill and return predictability – and in Table 1 panel B, I eliminate stocks below the 20th percentile of NYSE market capitalization. Including small caps results in higher point estimates and standard errors.[16] For other hypotheses, where I consider the joint behavior of asset prices and hedge fund trading rather than focusing solely on hedge fund skill, I include all stocks with available data to maximize the power of my tests.[17]

Table 1 panel B displays summary statistics of the volume consumed (aggregation method 1) decile portfolios. In the top-decile portfolio, volume consumed averages 9.3%: hedge funds trade very aggressively in these stocks (for an illustrative example, see Appendix B.1).[18] As one moves from lower to higher volume-consumed portfolios, the value of trades increases while total stock-level dollar volume (and market cap) decreases – neither effect operates in isolation. The top decile of purchases by volume consumed represents $12 billion of quarterly trades on average, or roughly $30 billion by the end of the sample.

4 Volume consumed: price impact and information

4.1 Contemporaneous performance – Hypothesis (1)

Hypothesis (1) suggests that trades generate price impact as they occur – regardless of information content – and that the magnitude of price impact should be a function of volume. In other words, high hedge fund volume consumed in quarter t should be associated with a high return in quarter t. As a first step, I test and confirm this prediction. My subsequent tests present evidence that these trades contain information.

Table 2 shows that stocks with high volume consumed in quarter t have very strong returns in quarter t. The table estimates Fama-MacBeth regressions with returns in quarter t as the dependent variable and the volume consumed quintile during quarter t – 1-5 for stocks with hedge fund trades, with 5 representing the highest volume consumed, and 0 for stocks with no hedge fund trades – as the explanatory variable. The coefficient on the volume consumed quintile is highly significant for all three methods of constructing volume consumed.[19] In column 1, for example – which focuses on stocks with non-zero hedge fund volume consumed – the coefficient estimate of 0.28% (t-stat 8.90) implies that stocks in the top quintile of volume consumed outperform stocks in the bottom quintile by 1.1% per month

[16]In the other direction, in Appendix B.2 I demonstrate that limiting the sample to stocks above the NYSE median market cap or NYSE median dollar volume still produces highly significant results.

[17]In unreported results, I find that limiting the sample throughout to stocks above the 20th percentile of NYSE market cap produces very similar results.

[18]Total volume figures include intra-quarter round-trip trades, which I do not observe from 13F filings. The trades that I identify may represent a greater proportion of non-round-trip volume.

[19]Table E.2 illustrates this result using decile portfolios.

on a characteristic-adjusted basis. As this performance is contemporaneous with hedge fund purchases, it is not a "tradeable" strategy from the perspective of an external observer of quarterly hedge fund holdings, even one who is not subject to the 45-day 13F filing delay. One would need to trade *before* a fund's first trade in order to capture this outperformance. This performance is also subject to a degree of survivorship bias: a stock must exist at the end of quarter t in order to appear in a 13F filing. The cross-sectional estimates based on stocks with non-zero hedge fund activity (volume consumed quintiles 1-5) should mitigate this issue. Still, these estimates may provide an upper bound for the price impact of trades.

The empirical microstructure literature provides robust evidence that large trades generate short-horizon price impact. If a portion of this price impact is permanent, then we would expect large quarterly trades to generate detectable price impact, too.[20]

The active management literature notes that some fund managers buy high momentum stocks (Grinblatt, Titman, Wermers (1995)) or add past winners to their portfolios (window dressing, Lakonishok, Shleifer, Thaler, and Vishny (1991)). While these effects may explain a portion of the contemporaneous outperformance I identify, it is not clear why they should be so much stronger for high volume-consumed positions than for low volume-consumed positions. If trades do not cause price impact, then managers should size trades based on absolute dollars (discussed below) or relative to their own portfolios (see Section 7.2).

Hedge funds are not naively purchasing stocks with high past risk-adjusted returns. In the Kyle model, future returns are a martingale. Some stocks are underpriced – those that had high past returns due to informed purchases – while other stocks are overpriced – those stocks that had high past returns due to noise trader purchases. In Table 2 column 6, I regress characteristic-adjusted quarter t+1 returns on a stock's quintile of quarter t characteristic-adjusted returns.[21] There is no evidence for characteristic-adjusted return continuation: the coefficient on the quarter t return quintile is insignificant. Without knowing hedge fund volume, past performance does not predict future characteristic-adjusted returns at these horizons. In contrast, as I show in Section 4.2, stocks with high hedge fund volume consumed in quarter t also have high characteristic-adjusted returns in quarter t+1.

Funds' largest purchases by dollar value (without reference to volume) also fail to display strong outperformance. To illustrate this point, I sort stocks into quintiles by the dollar value of trades calculated using the quarter's opening prices: I sort by $valoftrade_{s,t}^{open} = shares\,traded_{s,t} * P_{s,t-1}$, where $P_{s,t-1}$ is the price of stock s at the end of quarter t-1.[22] I use opening prices because closing prices have mechanical look-ahead bias, since $P_{s,t}$ is a function

[20]See Section 6 for estimates of the price impact function and comparisons to existing estimates.

[21]Results are similar if I use the quintile of past market-adjusted returns.

[22]This construction is parallel to aggregation method 1 for volume consumed. Results are similar using a construction parallel to aggregation method 2 or 3.

of returns during t. Note that volume consumed does not have mechanical lookahead bias, since neither shares traded during quarter t nor volume during quarter $t-1$ is clearly a function of returns during quarter t. In column 7, I find that in a regression of quarter t characteristic-adjusted returns on a stock's quintile of quarter t opening value of trades, the estimated coefficient is negative. Though using opening prices is an imperfect proxy, sorting by the opening dollar value of purchases suggests contrarian behavior. Regardless, in the absence of price impact, it seems difficult to explain why fund managers would window dress based on volume consumed rather than dollar values.

The existing literature that links trades to momentum finds long-horizon return reversals. Lou (2012), for example, shows that flow-induced mutual fund trades outperform contemporaneously, but that return reverses within three years. As I show in Section 4.5, the returns of high volume-consumed hedge fund trades do not reverse. Under a causal interpretation, flow-induced mutual fund trading pushes prices away from fundamental value. Hedge fund trading, in contrast, pushes prices towards fundamental value. Under a non-causal interpretation, hedge funds consume a large amount of volume in the subset of high contemporaneous return stocks that will not feature future long-term return reversals.

4.2 Predicting future returns – Hypothesis (2)

Willingness to move prices reveals information: the cross-section of hedge fund volume consumed is a powerful predictor of future stock returns (Hypothesis (2)). This test most closely corresponds to standard tests of investment skill and return predictability.

Table 3 panel A displays the results of Fama-MacBeth regressions that predict returns during quarter t+1 using the quintile of volume consumed in quarter t. The associated coefficient is positive and highly significant for all methods of constructing volume consumed.

Panel B presents decile portfolio returns, and removes stocks in the bottom quintile of market cap to emphasize that hedge funds are not only identifying mispricings in microcaps. Using aggregation method 1, the top-decile portfolio outperforms the lowest decile portfolio by 0.74% (0.55%) a month – 9.3% (6.8%) annualized – on a market-adjusted (characteristic-adjusted) basis, with a t-stat of 5.36 (4.56). On its own, the top-decile portfolio outperforms by 0.70% (0.47%) on a market-adjusted (characteristic-adjusted) basis, with a t-stat of 4.84 (5.76). The Kyle model interprets this outperformance as the result of a combination of continued trading (and its associated price impact) and the release of information.

These results suggest that large hedge fund purchases are highly informative for future returns. This finding is quite broad: the top five decile portfolios (using aggregation method 1) – half of hedge fund purchases – generate statistically significant outperformance. By dollar value, these trades represent 79% of hedge fund equity purchases. Characteristic-adjusted performance monotonically increases as one moves from decile 6 to 10.

The statistical significance of my findings is strong, especially for a portfolio that is not composed of microcaps or heavily reliant on short sales (and thus subject to short-sale constraints). As I demonstrate in Section 7.2 and Appendix B.2, this finding is not subsumed by previously identified empirical effects. The predictive power of volume consumed for future returns is robust to alternative explanations including downward sloping demand, heterogenous beliefs, fund activism, fund concentration, hot hands, asset flows, or simple sorts by volume, volatility, bid-ask spreads, PIN, or Amihud ratios. In a four-factor regression, the top-decile portfolio's loadings on MKT, SMB, HML, and UMD are 0.98, 0.58, 0.27, and -0.06, respectively (the four-factor alpha is 0.56% per month, with a t-statistic of 5.28). In a six-factor regression that adds the Pastor-Stambaugh value-weighted liquidity factor and the Sadka liquidity factor, the alpha is 0.51% per month, with a t-statistic of 5.39. Hedge funds are not simply taking momentum or liquidity risk.[23]

Sales are uninformative, as illustrated by the results for aggregation method 3 in panel B.[24] The bottom decile of positions by volume consumed (the largest sales) demonstrates slightly positive, though insignificant, future monthly outperformance of 0.22% (0.08%) on a market-adjusted (characteristic-adjusted) basis, with a t-statistic of 1.51 (1.08). This finding may reflect liquidity driven reversals. Part of the price impact of large trades is temporary, and dissipates even in the absence of information. Temporary price impact may reduce the outperformance of high volume-consumed positions.

To put these results into the context of my full sample, the portfolio that combines by value and holds all hedge fund positions generates market-adjusted (characteristic-adjusted) performance of 0.23% (0.12%) per month, with a t-statistic of 3.94 (3.00).[25,26] On an equal-weighted basis, the positions that hedge funds own generate market-adjusted (characteristic-adjusted) performance of 0.17% (0.07%) per month, with a t-statistic of 1.39 (2.09). My analysis focuses on purchases. The portfolio of all hedge fund purchases in the sample (aggregation method 1), on an equal-weighted basis, generates market-adjusted (characteristic-adjusted) performance of 0.18% (0.06%) per month with a t-statistic of 1.39 (2.03). Weighting by trade size, hedge fund purchases generate market-adjusted (characteristic-adjusted) outperformance of 0.32% (0.21%) per month, with a t-statistic of 3.65 (3.81).

[23]Industry-adjusting returns using the 48 equal-weighted Fama-French portfolios produces similar results.

[24]Di Mascio, Lines, and Naik (2015) also find that sales are uninformative in their sample.

[25]My analysis in this paragraph also eliminates stocks below the 20th percentile of NYSE-market cap. Including all stocks has a minimal impact on value-weighted figures. Stocks in the bottom quintile of market cap comprise less than 5% of the aggregate hedge fund portfolio by value.

[26]For comparability, Griffin and Xu (2009) find annualized market-adjusted (characteristic-adjusted) value-weighted outperformance of 0.21% (0.18%) in their hedge fund sample from 1986-2004. They include all hedge funds, without an attempt to focus on funds with informative 13F portfolios.

4.3 Predicting future trading – Hypothesis (3)

Hypothesis (3) says that if the private information has not been revealed, then the informed trader will continue to buy the most in quarter t+1 of positions with the highest volume consumed in quarter t. She spreads out her large trades across time to minimize price impact. If she does not cause price impact, then this behavior is systematically suboptimal.

Table 4 uses Fama-MacBeth regressions to show that high volume consumed this quarter predicts high volume consumed next quarter.[27] The coefficient estimate from column 1 suggests that volume consumed in quarter t+1 is 3.0% higher for stocks in the top quintile of volume consumed during quarter t than it is for stocks in the bottom quintile of volume consumed during quarter t. The coefficient is highly significant, with a t-statistic of 30.21.

This calculation uses aggregation method 1, which sums only purchases, to construct volume consumed. Examining specifically the sum of purchases next quarter is essentially proxying for information *not* having been released – in the sense of the model – at the start of quarter t+1. If information is released late, the informed trader will buy more in both quarter t and quarter t+1 of a position with a greater initial mispricing. This finding is by no means empirically obvious. For example, suppose that any time a hedge fund manager decides to build a new position, she decides to take up 5% of the stock's volume over the 90 days following her first trade. In that case, high volume consumed in quarter t would predict low volume consumed in quarter t+1, and vice versa.[28]

My empirical results imply that during quarter t+1, funds continue to buy more of the positions that do the best in quarter t+1 (positions with high volume consumed during quarter t have higher returns in quarter t+1). If funds do not cause price impact, then they are systematically leaving money on the table by not building these positions even faster (i.e., they should buy more during quarter t instead). If funds do cause price impact, it is optimal for them to spread out large purchases across time.

Column 5 displays results using aggregation method 3 to construct volume consumed in both quarters t and t+1. This method of constructing volume consumed may be of particular interest here because it nets purchases and sales. It shows that stocks subject to large net hedge fund sales in quarter t are followed by relatively large net sales in quarter t+1, while stocks subject to large net hedge fund purchases in quarter t are followed by relatively large net purchases in quarter t+1. For example, stocks in the bottom decile of volume consumed in quarter t (largest sales) average -1.21% of volume consumed in quarter t+1, while stocks in the top decile of volume consumed in quarter t (largest purchases) average 0.98% of volume

[27]Table E.3 illustrates this result using decile portfolios.

[28]In this case, if you see volume consumed in t of 4%, then you would see volume consumed in $t+1$ of 1%. If you see volume consumed in t of 1%, then you would see volume consumed in $t+1$ of 4%.

consumed in quarter t+1.

Hedge funds spread more of their largest trades across multiple quarters, as predicted by the model.[29]

4.4 Informed trading reduces the impact of a positive earnings surprise – Hypothesis (4)

Hypothesis (4) suggests that informed trades incorporate information into asset prices prior to the information's public release. I confirm this prediction using earnings announcements. As a precursor to this finding, I demonstrate that hedge funds unconditionally predict earnings returns before controlling for the magnitude of the information contained in the announcement. This result is evidence that funds predict company fundamentals.

4.4.1 SUE framework

To examine this hypothesis, I use the ex-post observable standardized unexpected earnings surprise (SUE) in quarter t+1 as a proxy for the initial mispricing. I focus on earnings announcement days because they contain substantially more information than other trading days. I treat the earnings release date as an "annnouncement" date in the model.

I study (weakly) positive earnings surprises because my data is informative regarding the information content of purchases. In this context, the theory has a clear implication for positive news: informed purchases should incorporate some of that information into prices prior to its release.[30] The implication for negative earnings surprises is not as straightforward. If hedge funds are informed, then for stocks with negative earnings surprises, funds are more likely to have been buying on the premise of something other than earnings news. On the firm side, some executives may gradually release negative news over time (Cohen, Lou, Malloy (2014)), which further complicates negative earnings releases.

For the same SUE (initial mispricing), stocks subject to more informed trading in quarter t should react less to the release of that fixed amount of information in quarter t+1. In a regression with the earnings return on the left hand side, the effective coefficient on SUE should be smaller for stocks with higher volume consumed the previous quarter. I therefore interact a stock's volume consumed quintile with its SUE and hypothesize that $\beta < 0$ in:

$$\begin{aligned} earningsreturn_{s,t+1} \quad = \quad & \beta\,(VCQ_{s,t} * SUE_{s,t+1}) + \alpha_1\,VCQ_{s,t} + \alpha_2\,SUE_{s,t+1} \\ & + \gamma\,controls_{s,t} + \nu_{s,t} \end{aligned} \tag{5}$$

[29] Di Mascio, Lines, and Naik (2015) also find that institutions in a proprietary dataset trade in the same direction in the same stocks over multiple quarters. They do not normalize trades by volume, however.

[30] I focus on positive earnings surprises rather than positive earnings announcement returns because the premise of this test is that the earnings surprise is not causally affected by hedge fund trading activity. If hedge fund purchases have price impact, the announcement return *will* be causally affected by trading.

In the data, to ensure that earnings and analyst forecasts reflect the same time period that hedge funds are trading over, I include only companies with calendar quarter-end fiscal periods and only use analyst forecasts made during calendar quarter t. The earnings return is measured as the return over the three trading-day window centered around the Compustat earnings announcement date, using characteristic-adjusted daily returns. Standardized unexpected earnings, $SUE_{s,t+1}$, is measured as $\frac{earnings_{s,t+1} - median\ analyst\ forecast_{s,t}}{P_{s,t}}$, as in Baker, Litov, Wachter, and Wurgler (2010). Additional data details are in Appendix D.1.

Table 5 panel A provides context, using Fama-Macbeth regressions. I employ market cap, dollar volume, book-to-market, and institutional ownership as control variables. The first column is of interest in its own right: it demonstrates that volume consumed in quarter t predicts returns during the earnings announcement window in quarter t+1. This finding is evidence that funds predict fundamental information in the stocks they purchase heavily. Roughly 25% of the total characteristic-adjusted outperformance of the top-decile portfolio is realized during the earnings window, even though the earnings window encompasses fewer than 5% of the trading days in the average quarter.[31] Volume consumed also predicts returns on other days in the quarter (column 2). I then shift to focus on stocks with (weakly) positive earnings surprises. In column 3, I show that volume consumed continues to predict earnings returns in this sample. In column 4, I find that SUE strongly predicts the earnings returns of stocks in this sample. In column 5, I find that SUE partially displaces hedge fund volume consumed in a horserace: the coefficient on the volume consumed quintile drops from 0.06% in column 3 to 0.03% in column 5, and loses statistical significance. SUE is measured contemporaneously, whereas volume consumed has to predict the earnings return using information from the previous quarter. The horserace suggests that hedge funds are partially predicting next quarter's earnings return, but that SUE is a more accurate measure of that information. Supporting this interpretation, column 6 shows that the quintile of volume consumed positively predicts the magnitude of positive SUE. These findings reinforce my treatment of SUE as a proxy for information in this sample.

4.4.2 Reaction to positive SUE following informed purchases

Table 5 panel B shows that stocks with more informed trading appear to be more efficiently priced *prior* to the public announcement of positive earnings surprises. The point estimate suggests that 28% of earnings information is incorporated prior to its release in stocks in the top quintile of volume consumed.

Fama-Macbeth regressions estimate a significantly negative coefficient on the interaction

[31] Baker, Litov, Wachter, and Wurgler (2010) show that mutual fund trades also have some ability to predict earnings returns in their sample.

of the volume consumed quintile and SUE (equation (5)). That is, the *effective* coefficient on SUE – the coefficient on SUE plus the coefficient on the interaction term multiplied by a stock's volume consumed quintile, $\alpha_2 + \beta VCQ_{s,t}$ – declines for stocks in higher volume consumed quintiles.[32] The point estimate on the interaction term ranges from -0.13% to -0.29%, with a t-statistic between -2.11 and -3.63.

The associated economic significance is substantial. The estimate on the full sample (column 1) implies that moving from stocks with no hedge fund volume consumed to stocks in the top quintile of volume consumed reduces the effective coefficient on SUE from 2.52% to 1.82%, a decline of 27.8%.

Columns 2-6 employ different regression specifications and subsamples. In those columns, I also interact control variables with SUE. My findings continue to hold after eliminating stocks in the bottom quintile of market cap or stocks with no hedge fund volume consumed, in pooled regressions using double-clustered standard errors (by firm and quarter), or including firm fixed effects and clustering standard errors by quarter.

If hedge funds are aware of differences in stocks' reactions to SUE – even after controlling for observables such as volume – then funds could drive my results by endogenously choosing to purchase stocks that react less to positive earnings. However, note the implication of such endogeneity for fund profits. If hedge funds can predict SUE *and* how responsive a stock will be to that SUE, then funds should seek to purchase stocks with high SUE and *high responsiveness* to that SUE in order to maximize fund returns. Instead, funds consume the most volume in stocks with high SUE but low responsiveness, a situation they should seek to avoid if they can. Clearly funds would *want* to trade without generating price impact.

Measurement error is another potential concern. Perhaps hedge funds invest the most in stocks for which SUE is simply noisier. Assuming classical measurement error in SUE, this would cause β to be biased negatively away from zero (the effective coefficient on SUE would decline as one moved from stocks without hedge fund activity to stocks with high hedge fund activity). In that case, however, SUE should also have less explanatory power for the earnings returns of high volume-consumed stocks. I find the opposite to be true. I run equation (5) separately by volume consumed quintile, and find that the r-squared is highest for stocks in the top quintile of volume consumed.[33,34]

[32] Table E.4 illustrates this result by estimating the coefficient on SUE separately for three groups – (1) stocks with no hedge fund volume consumed or in the bottom quintile of volume consumed; (2) stocks in the middle three quintiles of volume consumed; and (3) stocks in the top quintile of volume consumed.

[33] I run six separate Fama-MacBeth regressions of $earningsreturn_{s,t+1} = \alpha_2\, SUE_{s,t+1} + \gamma\, controls_{s,t} + \nu_{s,t}$. For stocks in the top quintile of volume consumed, the r-squared is 12.0%. The r-squared ranges from 6.9% to 8.7% for stocks in other volume-consumed quintiles (or stocks without any volume consumed).

[34] A related concern is that perhaps the effective coefficient on SUE is simply smaller for high SUE. Of course, that *is* an equilibrium outcome of the model: if hedge funds predict high SUE and reduce the reaction of the stocks they trade to a given SUE, then on average stocks with high SUE will react less per unit of SUE

This test associates the reduced reaction of stocks in high volume-consumed quintiles to hedge fund activity. Hedge fund activity may be a sufficient statistic for the trading activity of all arbitrageurs prior to the release of earnings. However, 94% of earnings reports are released prior to the public 13F filing date.[35] Nevertheless, there could be some market participants who reach the same investment conclusions as and trade contemporaneously with (and in the same direction as) hedge funds. The greater information content of prices prior to the release of earnings remains of interest even under a sufficient statistic interpretation. The price impact associated with a specific amount of trading would be reduced, however.

These results suggest that hedge fund activity substantially reduces the reaction of stock prices to positive earnings announcement surprises.

4.4.3 Reaction to negative SUE following informed purchases

For completeness, Table 5 panel C reports the reaction of stocks to negative SUE following different levels of hedge fund purchase activity during the previous quarter.

As discussed above, the asymmetry of the informativeness of purchases relative to sales in a long portfolio leads me to focus on positive earnings surprises. In contrast, the results from negative earnings surprises do not display a clear trend. The coefficient on the interaction term is small in magnitude and is not significantly different from zero at a 5% level in any of the specifications. The coefficent on the volume-consumed quintile – which picks up non-earnings news – remains significantly greater than zero. Comparing column 1 across panels B and C, the coefficient on the volume consumed quintile is nearly twice as large for stocks with negative SUE as for stocks with positive SUE. This result suggests that hedge funds are more likely to be trading on non-earnings information in stocks with negative SUE. Those trades may not be simply "mistakes." Furthermore, the coefficient on negative SUE is substantially smaller than the coefficient on positive SUE. Comparing column 1 across panels B and C, the coefficient on positive SUE is 2.52%, while the coefficient on negative SUE is about one-sixth that. The information content of negative earnings surprises differs from that of positive earnings surprises.

4.5 Long-horizon cumulative returns – Hypothesis (5)

For Hypothesis (5), I test whether the outperformance that I identify persists over long horizons. The Kyle model considers permanent price impact and information. The returns associated with hedge fund trading should not revert. This test further rules out the possi-

because hedge funds trade more of those stocks. Nevertheless, first note that in Table 5 panel A the linear relationship between positive SUE and the earnings return is strong (SUE has a t-stat of 10.78). Second, the coefficient on the interaction term between the volume consumed quintile and SUE remains negative and significant (t-stat -1.86) after removing all stocks with top-decile positive SUE.

[35]Dropping observations of earnings released after the 13F date does not impact my results.

blity that hedge funds merely predict the future price impact of their own or others' trades.

Figure 2 shows that the cumulative buy and hold outperformance of the top volume-consumed decile portfolio (aggregation method 1) remains significantly positive, relative to the bottom volume-consumed decile portfolio, for 2-4 years. Cumulative performance is calculated by forming calendar-time portfolios.[36] These portfolios go long the top decile of positions by volume consumed in quarter t and short the bottom decile of positions by volume consumed in quarter t. Panels A and B use only future performance figures: they look at performance during quarter t+1 and beyond. Cumulative future outperformance is significantly positive for roughly 2-3 years, reaching about 9% (4%) on a market-adjusted (characteristic-adjusted) basis at a five-year horizon.

Panels C and D give credit to the contemporaneous performance that accompanies a manager building her positions in period t: they look at performance beginning at the start of quarter t. In this case, outperformance remains significantly positive for 4-5 years, reaching about 16% (10%) on a market-adjusted (characteristic-adjusted) basis at a five-year horizon. Section 4.1 makes the case that much of the contemporaneous price movement in these figures reflects price impact.

While price impact may affect price movements over multiple quarters, it is unlikely to cause price movements that are fully persistent at a five-year horizon. At that horizon, information is more likely to be the primary determinant of the cross-section of asset prices, at least on average. Hedge fund trades are not persistent over five years, for example (as I show in Appendix B.2). In contrast to these results, large mutual fund trades are associated with price movements that tend to revert – as in Section 5.1 and in the flow-driven trading literature – within three years at the longest.

In this sense, hedge fund trades are associated with and may partially drive the movement of asset prices towards their long-run fundamental values.

5 Mutual funds

In the Kyle model, there are two types of active traders: informed traders and uninformed (noise) traders. Mutual funds could potentially fall into either group. In either case, the large trades of mutual funds should be associated with strong contemporaneous returns – all trades generate price impact. In some sense, this test is an out-of-sample test of one of the fundamental premises of the model. If mutual funds are informed, then their large trades

[36] These portfolios are constructed as overlapping portfolios, as in Jegadeesh and Titman (1993) or Coval and Stafford (2007). In particular, these results represent the return to a strategy that purchases (shorts) the top (bottom) decile of volume-consumed positions each quarter and holds them for the relevant time horizon. In any given calendar month, the portfolio is then equal weighted across the long-short portfolios that were formed at each relevant formation date. For horizon k returns, in calendar quarter $t+1$ the portfolio return is the equal-weighted average of the quarter $t+1$ returns of the k long-short portfolios that were formed at the end of quarters $t-k+1$, ..., t.

relative to volume should also predict future returns, and those returns should not revert.

I find that large mutual fund trades comove with contemporaneous returns but are uninformative for future returns (at least compared to hedge fund trades). However, when I limit the sample to subsets of mutual funds that previous research suggests may be differentially skilled, I find that volume consumed does in fact positively predict future returns.

5.1 Mutual fund volume consumed and returns

As the model predicts, large mutual fund purchases are associated with high contemporaneous returns. However, I find only weak evidence that mutual fund volume consumed predicts returns during the following quarter. Furthermore, these returns revert at multi-year horizons, which should only be true for uninformed purchases.

I examine a sample of mutual funds that is comparable to my hedge fund sample. I limit my tests to mutual funds with between 10 and 150 positions, and at least $50 million in assets. I include only mutual fund filings that occur at calendar quarter ends (and I only calculate trades from filings at contiguous quarter ends). To remove index and target date funds, I eliminate funds with "index" or its variations, or future dates (2025, 2030, etc.), in their fund names. I eliminate international, municipal bonds, bonds and preferred, and metals funds (IOC codes 1, 5, 6, and 8).

The literature documents that extreme mutual fund flows drive price dislocations. To differentiate my findings, I eliminate trades by funds in the top and bottom deciles of flows (flows are defined as in Coval and Stafford (2007); see Appendix D.2 for details). I construct volume consumed by aggregation method 1 (purchases aggregated at the stock level).[37]

Table 6 displays the performance of mutual fund trades. Panel A focuses on the broad mutual fund sample. The model suggests *all* trades should generate price impact as they occur, regardless of information content. Panel A columns 1 and 2 confirm this prediction. In Fama-MacBeth regressions with returns in quarter t as the dependent variable and the mutual fund volume consumed quintile during quarter t as the explanatory variable, the coefficient on the volume consumed quintile is positive and highly significant.[38]

Mutual funds trades do not significantly predict future returns, however. Panel A columns 3 and 4 present results from regressing future returns on the mutual fund volume-consumed quintile. The associated coefficient is insignificant.

Large hedge fund trades relative to volume outperform large mutual fund trades relative to volume. To make this comparison, I form a long-short portfolio that is long the top decile of hedge fund trades by volume consumed in quarter t, and short the top decile of mutual fund trades by volume consumed in quarter t. This portfolio returns 0.40%

[37]Results are similar using aggregation methods 2 and 3 to construct volume consumed.

[38]Table E.5 illustrates the results of this section using decile portfolios.

(0.31%) per month during quarter t+1, with a t-statistic of 3.92 (3.06) on a market-adjusted (characteristic-adjusted) basis. The top hedge fund trades substantially outperform their mutual fund counterparts. However, mutual fund trades in the top decile take up about 7.5% of volume, on average, compared to 9.3% for top-decile hedge fund trades. Yet even the 9th decile of hedge fund purchases – associated with 4.5% of volume consumed – performs significantly better than the top decile of mutual fund purchases. The associated long-short portfolio generates performance of 0.23% (0.21%) per month, with a t-statistic of 2.21 (1.94), on a market-adjusted (characteristic-adjusted) basis.

The cumulative returns of stocks heavily bought by hedge funds are significantly positive at multi-year horizons. In Figure 3, I display comparable returns for mutual fund trades (the figure also displays hedge fund results for reference). The modest outperformance of high volume-consumed mutual fund trades reverses. Using only future returns in panel A, long-horizon cumulative returns become negative after about a year. These figures have large standard errors and are not statistically significantly different from zero. They *are*, however, statistically different from the long-horizon returns of hedge fund trades. Even after including the strong contemporaneous performance of mutual fund trades in panel B, returns revert within two years. These results are consistent with mutual funds as uninformed traders. In the model, noise trades have price impact, but that price movement reverses as the release of information pushes prices back towards fundamental value.

This finding is similar to the existing literature on flow-induced mutual fund trading, i.e., Coval and Stafford (2007), Frazzini and Lamont (2008), and Lou (2012). My results hold even after removing the purchases of mutual funds subject to extreme flows. Khan, Kogan, and Serafeim (2012) also find reversals following general mutual fund buying pressure.

5.2 Mutual fund skilled subsets

The literature provides evidence that we may be able to identify skilled subsets of mutual funds. If skilled funds reveal information via volume, and we can identify skilled funds, then the volume consumed by their trades should predict future performance. I confirm this hypothesis using funds in the top quintile of return gap and funds with above-median active share as plausibly skilled funds.

I construct funds' return gaps following Kacperczyk, Sialm, and Zheng (2008) (see Appendix D.3 for details). I lag the measure of return gap by three months: I examine the performance during quarter t+1 of trades during quarter t by funds in the top (or bottom) quintile of return gap from the end of quarter t-5 to the end of quarter t-1. I introduce this extra quarter lag because a fund that consumes a large amount of volume in quarter t could potentially cause its return gap to increase by generating price impact. I seek to differentiate my findings from this possibility. I take data on active share from Antti Petajisto's website.

Table 6 panel B illustrates my findings. In the skilled subsets, the coefficient on the volume-consumed quintile is significant at a 10% level in a regression with future characteristic-adjusted returns as the dependent variable (columns 1 and 3). In contrast, the volume consumed by funds in the bottom quintile of return gap or funds with below-median active share is not informative (columns 2 and 4). Even the trades of the plausibly skilled mutual funds are substantially less informative than hedge fund trades, however. The coefficient on the volume-consumed quintile is more than twice as large for hedge fund trades (Table 3).

6 Quantifying the price impact function

6.1 Reduced-form and structural approach

I have presented evidence that the comovement of hedge fund trades and asset prices is consistent with the Kyle model. Up to this point, I have primarily utilized portfolio sorts. The model has more precise quantitative implications for the permanent price impact function, however: price impact is linear in net order flow, with coefficient λ. This parameter is of economic interest because it determines how much information a given amount of informed trading incorporates into prices and affects how quickly returns to scale diminish in asset management.

I estimate λ by two approaches. In the first, I directly estimate the reduced form: I regress contemporaneous returns on informed trading. In the second, I impose the full set of equilibrium constraints and structurally estimate the model from Section 2 via maximum likelihood. Hedge fund optimization implies that funds should trade in a certain manner given their knowledge of the price impact function. The structural model applies this intuition to observations of trading to infer what fund managers believe λ to be. The model supplements the reduced-form equation with this information to estimate λ. However, the model requires additional assumptions in order to apply this economic structure to the data.

I find that reduced-form estimates imply that consuming 1% of quarterly volume generates 0.3% of price impact, while comparable structural figures range from 0.3% to 0.5%. As evidence that the model provides a reasonable quantitative description of the data, I find that model-simulated moments of trading and returns are close to the empirical moments.

These estimates offer a complementary perspective on Kyle's λ relative to existing microstructure estimates. For reduced-form estimation, the microstructure literature associates intraday returns with trades. The Kyle model is specifically a model of permanent price impact, however. By employing quarterly returns, I allow time for temporary price impact to dissipate. Microstructure procedures must introduce additional parameters to attempt to control for temporary price impact. On the other hand, high-frequency microstructure estimates can more easily separate return chasing from price impact.

For structural estimation, the microstructure literature often assumes that information episodes last a single day. However, there is evidence that trades are coordinated over longer periods, as Kumar and Lee (2006) and Collin-Dufresne and Fos (2015) show using proxies for noise and informed trades, respectively. I use two-quarter information episodes and quarterly time periods in my estimation. While this allows trading to be coordinated across longer time periods, quarters are unlikely to be precisely the correct unit of observation, either. On a separate note, microstructure models often indirectly identify informed trades from anonymous trading data. I directly employ data on plausibly informed hedge fund purchases. As I only observe long portfolios, however, my set of informed trades is censored.

6.2 Implementation

I assume that hedge fund volume consumed proxies for informed trading and that characteristic-adjusted returns proxy for price movements.

The reduced-form estimate links contemporaneous returns and trading without any additional structure on trading. In contrast, the model assumes that hedge funds are informed and *imposes* that they optimally scale trades given a signal of future returns and knowledge of the price impact function. Mathematically, the reduced form fits equations (3) and (4). The model jointly fits empirical moments (1)-(4), subject to equilibrium constraints (6)-(10).

The model considers permanent price impact. To the extent that trades incur other costs that increase the marginal cost of trading – and thus alter the informed trader's first-order condition – the model may overstate how aggressively funds should trade given information.

I build the structural likelihood function in Appendix C.1. I note here the key assumptions, which link the model to the data in a simple and direct manner. These assumptions are strong, however. One could draw on evidence from alternative settings or expand the model to supplement these assumptions at the cost of reduced transparency.

First, because I am only able to proxy for informed purchases, my data is censored. Maximum likelihood allows me to explicitly model censoring in the likelihood function. I use volume consumed of 0.1% as the censoring cutoff for both of my estimation procedures.

Second, I do not directly observe the revelation of information. I therefore do not use the sharp returns implied by the "revelation event" in the model to estimate parameters. I assume that continued informed buying in the second quarter (above the point of censoring) implies that information has not yet been released.

Third, I structure information episodes as exactly two quarters long. I study non-overlapping two-quarter intervals and assume that a new piece of information is simultaneously generated for each stock at the beginning of every interval.

Fourth, I estimate a single λ coefficient. Differences in the information environment across stocks or time could potentially drive corresponding variation in λ.[39]

Based on these assumptions, I estimate two parameters prior to maximizing the full likelihood function. First, I estimate the probability that information is released late, π, as the proportion of times that hedge fund purchases in quarter 1 are followed by further purchases in quarter 2. Second, I estimate the variance of information, σ_ϵ^2, as the variance of two-quarter returns.

6.3 Simulated data

To build intuition for how my estimation procedures differ, I consider how reduced-form and structural estimates of λ_1 change when I vary trading and returns in simulated data.

Consider four equal-frequency sets of stocks. I describe the first two sets, and assume that the second two sets are mirror images of the first two (with exactly -1 times their trading and returns). The baseline scenario is as follows. In the first quarter, for the first set of stocks, trades are 5% of volume, and returns are normally distributed with mean 5% and standard deviation 3.5%. For the second set of stocks, trades are 2.5% of volume, and returns are normally distributed with mean 2.5% and standard deviation 3.5%. Trading and returns are drawn identically in the second quarter, except that I randomly set ($\pi =$) 50% of trades to zero to signify that information is revealed early half of the time. In this data, the reduced form estimates $\lambda_1 = 1.00$.[40] The structural estimate is a nearly identical $\lambda_1 = 0.99$. In each of the following scenarios, I vary a single moment relative to this baseline case.

First, I vary the noise in returns: I increase the standard deviation of returns from 3.5% to 7%. Noise in the dependent variable does not alter the reduced-form point estimate. The structural estimate, however, increases to $\lambda_1 = 1.31$. The model assumes that the informed trader observes a signal of future returns, and thus interprets variation in returns as information. The model infers that price impact must have increased (higher λ_1) if the informed agent's information increases but her trading does not.

Second, I add Gaussian white noise with a standard deviation of 1% to trades. In this case, noise in the independent variable drops the reduced-form estimate of λ_1 to 0.64. The structural estimate only drops to 0.95. The structural estimate does not drop as far because it imposes a relationship between informed trading and returns based on optimization. The model is not as quick to discard trades as noise.

Third, I shift mean returns: I add 1% to the returns of the first two sets of stocks. The reduced-form estimate absorbs the increase in mean returns in its constant term. The structural estimate of λ_1, however, increases to 1.14. In the model, returns in the first period

[39] Normalizing trades by total volume should reduce concerns regarding differences in noise trading. Many microstructure estimates that focus on cross-sectional variation in λ use unadjusted trade size (in shares).

[40] Between the two sets of uncensored stocks, $\frac{average\ \Delta returns}{average\ \Delta trades} = \frac{2.5\%}{2.5\%} = 1$.

are due to trading. If higher average returns are associated with the same amount of trading, the model infers that a given trade must move prices by more (higher λ_1).

Fourth, I add 1% to all non-zero trades in the first two sets of stocks. The reduced form still produces $\lambda_1 = 1.00$. The structural estimation now reduces λ_1 from 0.99 to 0.80. It assumes that informed trades are made solely based on expected future returns rather than for other motivations (which might produce a constant amount of trading across stocks). Given that returns have not changed, the model infers that if the informed trader is willing to trade more, then she must believe her trades generate less price impact (lower λ_1).

6.4 Estimates from hedge fund purchases

Table 7 presents the results of my structural and reduced-form estimations based on hedge fund purchases. In my discussion, I focus on structural estimates pertaining to the first quarter of each information episode; second quarter results, i.e. λ_2, are similar.[41] I structurally estimate the model using 10 or 100 volume-consumed sorted portfolios. I estimate the reduced form using 10 portfolios, 100 portfolios, or the full cross section.

Structural estimates of λ_1 range from 0.31 to 0.53, with standard errors based on time-clustered bootstraps of 0.03 to 0.04. Reduced-form estimates of λ_1 range from 0.29 to 0.31, with standard errors of 0.03. These estimates imply that consuming 1% of volume over a full quarter generates permanent price impact of 0.3%-0.5%. These estimates fall below linearly aggregated academic and practitioner estimates of comparable *total* price impact as a function of volume, estimated using short-horizon returns, of roughly 0.8%.[42]

The structural and reduced-form estimates differ because of the additional moments that the structural model considers. Those moments fit hedge fund trading as the outcome of an optimization process. Using 10 portfolios, estimating the model only on the additional moments – equations (1) and (2) – produces an estimate of λ_1 of 0.79. This result suggests that hedge funds internalize a higher cost of trading than is estimated by the reduced form. The existence of other costs of trading besides permanent price impact – which, as mentioned, is the focus of the model – may contribute to this difference. The final structural estimate of 0.32 combines the reduced-form estimate (0.29) and the alternative estimate (0.79), but puts greater weight on the reduced-form estimate. Intuitively, the model gleans more information from the reduced-form equations because they use data on both observed trading and observed returns. The additional moments only use data on observed trading, which they link to the unobserved (inferred) mispricing.

To get a sense of how well the model fits the data, a natural approach is to compare

[41] I estimate the reduced form using all quarters since I do not need to delineate information episodes.

[42] The simple calculations underlying this figure are detailed in Appendix C.3, and are based on estimates from Frazzini, Israel, and Moskowitz (2012), Collin-Dufresne and Fos (2015), Brennan and Subrahmanyam (1996), and Investment Technology Group.

the model-implied moments of trading and returns to their empirical counterparts. To do so, I simulate the model and censor the resulting "observations." As presented in Table 7, I find that the resulting moments are reasonably close to their empirical counterparts. Using 10 portfolios, for example, the model generates a mean (standard deviation) of informed trading of 4.6% (3.4%), compared to 3.8% (5.5%) in the data. The model generates a mean (standard deviation) of returns of 1.5% (2.6%), compared to 2.1% (3.5%) in the data.

Overall, the model does not generate as much noise as there is in the data. Intuitively, the model assumes that the informed agent trades *solely* based on information, and that all information is known in advance. Empirically, some trades are driven by other considerations and new information arrives over time, generating noise in trading and returns, respectively.

Averaging observations to more aggregated portfolios reduces the noise in returns to the extent such noise has a mean of zero.[43] Indeed, the empirical standard deviation of returns is substantially lower with 10 portfolios than with 100 portfolios. As explained in Section 6.3, the model interprets noise in returns as information. With less information but a similar amount of trading, the model reduces its estimate of λ_1. Consistent with this reasoning, the structural estimate of λ_1 based on 10 portfolios, 0.32, is lower than the 100-portfolio estimate of 0.53. The 10-portfolio estimate may be preferred if it is based on less noise.

Nevertheless, I also solve and estimate two extensions of the model (Appendix C.2) that reduce the estimate of λ_1 based on 100 portfolios. First, I assume the informed trader has constant absolute risk aversion, which reduces λ_1 from 0.53 to 0.43. Second, I assume both risk aversion and that an orthogonal public "new information" shock occurs every quarter,[44] which further reduces λ_1 to 0.31. The intuition is that these additions cause the informed agent to trade less, holding other parameters fixed. A risk averse agent trades less for a given amount of information, while new information shocks reduce the information available at the beginning of the episode (and trade size = β * available information). To trade the same amount as before, the informed trader must believe her trades generate less price impact.

Finally, it is worth noting that in the model, the informed trader sees information (ϵ) plus noise (η). At the release of information, the price moves to ϵ. Unfortunately, empirically the information structure is not as sharp. Without observing the information event, the structural estimate assigns all variation in the trader's signal to ϵ rather than η.

6.5 Aggregate price impact and information incorporation

If hedge fund trades are based on information, then the permanent price impact they generate captures the amount of information they incorporate into prices. If my estimates of

[43] Averaging does not reduce the variation in informed trading nearly as much because portfolios are formed after sorting by volume consumed.

[44] In this case, the variance of two-quarter returns equals $\sigma_\epsilon^2 + 2\sigma_{ni}^2$, with σ_{ni}^2 as the variance of new information. σ_ϵ^2 weakly decreases.

the permanent price impact function are valid, then I can calculate this amount. On average, hedge fund purchases take up at least 0.1% of volume – the point at which I censor my data – in 37% of stocks. Taking $\lambda = 0.30$ – roughly the bottom of my range of estimates – and multiplying this figure by volume consumed implies that hedge funds move the prices of the stocks they purchase by an average of 1.2% per quarter on a characteristic-adjusted basis.[45] In terms of the associated changes in market capitalizations, this means that hedge funds move market caps by an average of $14 million per stock-quarter, for a total across stocks of $36 billion per quarter. This figure is 0.5% of the total opening market cap of these stocks ($7.3 trillion), or 0.4% of the opening capitalization of the entire market ($9.3 trillion).

The standard deviation of quarterly returns averages 23.9% in the set of stocks in which hedge funds take up at least 0.1% of volume. Thus hedge fund purchases move prices by (1.2%/23.9%=) 5.0% of a one-standard deviation movement in returns. In this sample, the average r-squared of a Fama-Macbeth regression of returns on implied price impact is 0.9%.

Stock-level quarterly returns also reflect the creation of new information over time and the price impact of noise trades. Relative to the information available at the start of each two-quarter information episode, in the sense of the model, hedge fund trading may incorporate a greater share of information. For example, structural estimates imply that hedge funds incorporate 35-40% of available information in the first quarter of trading, which is close to the change in the coefficient on positive SUE between stocks in the top quintile of volume consumed and stocks with no volume consumed (28%, Section 4.4). This estimate likely represents an upper bound, since it assumes that hedge funds have an unbiased signal of all information that is available at the start of each information episode.

7 Additional evidence

7.1 Firm insider trades

There is strong evidence that the purchases of firm insiders are informative about the cross-section of future stock returns.[46] There is a clear information-based reason for these trades to outperform: a firm's executives are better informed about the future cash flows of the business than is a typical trader. As further evidence that hedge funds are informed about firm fundamentals, I find that hedge fund volume consumed positively covaries with insider trades. Yet insider trading does not subsume hedge fund volume consumed.

Table 8 examines how hedge fund and insider trades relate. I construct an indicator vari-

[45] These figures are calculated by forming the relevant quantity for each stock-quarter with volume consumed above the point of censoring, averaging across all such stocks each quarter, and then taking the time-series average across my sample. Volume consumed is constructed by aggregation method 1. All variables are winsorized at the 1%/99% levels. I convert nominal figures to 2012 equivalents using U.S. CPI.

[46] Among others, Jeng, Metrick, and Zeckhauser (2003) show that insider purchases earn abnormal returns.

able for insider purchases, set to 1 if firm insiders net purchase shares in stock s during quarter t (summing all Form 4 insider purchases (positive) and sales (negative)).[47] Column 1 first demonstrates that volume consumed forecasts the cross-section of characteristic-adjusted equity returns in this sample (which includes control variables but is not limited by analyst and earnings data). The coefficient on the volume-consumed quintile is positive and highly significant. Column 2 shows that insider purchases are also highly informative.

Hedge funds tend to buy alongside insiders. With the indicator for insider purchases in quarter t as the dependent variable, the coefficient on the volume-consumed quintile in quarter t is 0.0023, with a t-stat of 2.43 (column 3). This coefficient suggests that stocks in the top quintile of volume consumed are associated with a 1% higher probability of net insider purchases than stocks in the bottom quintile. The simple correlation between an indicator for net insider purchases and an indicator variable for a stock being in the top quintile (decile) of volume consumed is 0.016 (0.022).[48]

Column 4 shows that hedge funds do not appear to be merely following intra-quarter insider purchases. When volume consumed and insider purchases in quarter t are used together in a regression to predict returns in quarter t+1, both variables remain highly significant. There is essentially no change in the coefficient on the volume-consumed quintile.[49]

These results provide further evidence of the information content of hedge fund trades.

7.2 Idiosyncratic risk and portfolio weights (best ideas)

There are two leading intuitions for how a fund should trade based on private information: the fund should trade until it hits a limit of either (1) price impact or (2) idiosyncratic risk. In the first case, the fund trades until the next trade would moves prices so far that total profits would be reduced (the Kyle model). In the second case, the fund trades until it has assumed the maximum amount of idiosyncratic risk that the fund is willing to take on that position. In the classic limits-to-arbitrage story (Shleifer and Vishny (1997)), a fund that underperforms by a sufficient amount in the short run may be liquidated.

I illustrate that in my hedge fund sample, the first limit (price impact) seems to bind more closely than the second (idiosyncratic risk), in the sense that the former is statistically more informative for future returns. However, the measure of idiosyncratic risk that I employ may not be an effective proxy for the risk function of hedge fund managers. If position sizes are constrained by a different risk measure, perhaps positions with the highest risk weights by

[47] The literature typically uses an indicator variable when studying the information content of insider purchases. The trading activity of insiders is a small proportion of trading and is closely regulated, making the use of volume consumed inappropriate in that context.

[48] In unreported results, I do not find evidence for significant leads / lags in this relationship at the quarterly frequency of my data.

[49] The coefficient on the volume-consumed quintile is also essentially unchanged in a regression to predict future returns within the set of stocks that insiders do *not* net purchase.

that measure outperform. For example, Rhinesmith (2014) provides evidence that portfolio weights may be constrained by past losses in a stock.

I take my measure of idiosyncratic risk from Cohen, Polk, and Silli (2010, CPS). If an investment fund manager maximizes her portfolio's CAPM-adjusted Sharpe ratio, and ignoring trading frictions and price impact, her positions with the largest risk-adjusted portfolio weights should have the highest expected returns.[50] These are the fund's "best ideas."

CPS measure overweights relative to a stock's market cap weight, either out of the entire CRSP-value weighted index or out of the sum of the market capitalizations of all the stocks in a manager's portfolio. I use the former construction here, but unreported results are similar using the latter construction. CPS then multiply this overweight or underweight by a stock's idiosyncratic CAPM variance, which I measure using rolling windows of 36 months of returns.[51] In their mutual fund sample, CPS find that funds' top positions according to this measure significantly outperform.

Table 9 panel A compares volume consumed and best ideas. I analyze volume consumed by aggregation method 2 ($volconsumed_{s,f,t}$) for comparability to CPS, who analyze stock s, fund f, time t triplets. I form overlapping bins, and display the characteristic-adjusted future performance (and associated t-statistics) of the corresponding portfolios during quarter $t+1$. I form three groups by volume consumed: positions with no volume consumed or in the bottom quintile, positions in the middle three quintiles, and positions in the top quintile. I then independently group positions by their intra-fund best ideas ranking, as CPS do. I create three bins: positions with the top 3 values of best ideas for each manager, positions 4-10, and all other positions (11+).[52,53]

Positions in the top quintile of volume consumed outperform, regardless of their best ideas ranking. In contrast, stocks in the top group of best ideas significantly outperform only if they are also in the top quintile of volume consumed. Point estimates of abnormal returns are insignificantly positive for other positions ranked in the highest best ideas bin. In my hedge fund sample, volume consumed subsumes this measure of idiosyncratic risk.

7.3 13F filing dates

Hedge funds must file Form 13F within the first 45 days following the end of each calendar quarter. After filing Form 13F, the fund's holdings at the previous quarter end become

[50] Wurgler and Zhuravskaya (2002) motivate a similar approach. See Appendix E.6 for a brief discussion.

[51] That is, $bestideas_{s,f,t} = \sigma^2_{s,t,idCAPM}(w_{s,f,t} - w_{s,M,t})$, with $\sigma^2_{s,t,idCAPM}$ as stock s's idiosyncratic CAPM variance at quarter t, $w_{s,f,t}$ as fund f's portfolio weight in stock s at quarter t, and $w_{s,M,t} = \frac{mktcap_{s,t}}{\sum_{s=1}^{S} mktcap_{s,t}}$ as the weight of stock s in the value-weighted index at quarter t (the sum is over the set of all stocks S).

[52] In Table E.6, I show that results look similar using a finer partition.

[53] In unreported analysis, I find that results are similar if I pool positions across all managers before sorting by best ideas, instead of using a position's *intra*-manager ranking. Results are also similar if I measure idiosyncratic CAPM variance over 24-month windows or using 3-month windows of daily returns.

publicly observable. Table 9 panel B breaks down characteristic-adjusted performance of the extreme decile volume consumed (aggregation method 1) portfolios during quarter t+1 into three intervals: before the 13F-day window, the three trading-day window centered around the 13F filing date, and after the 13F-day window.

Outperformance for high volume-consumed positions remains significant both before (column 1) and after (column 3) the 13F-day window. Returns during the narrow three-trading-day 13F window (column 2) are insignificant. The differences between returns before and after the 13F filing window (column 4) are also insignificant.

These results suggest that a publicly-implementable long-short strategy could capture a portion of the outperformance that I identify. Historically, this strategy returned 0.58% monthly on a characteristic-adjusted basis during the second half of every calendar quarter.

Depending on the information structure, there could be a tension between the public implementability of this strategy and a strict interpretation of the model. If market participants know that hedge funds are informed, then prices should adjust as soon as hedge fund trades are publicly released (i.e., during the 13F-day window). In practice, however, an entire literature debates the information content of the trades of different investment managers. From the perspective of an econometrician in 2015 with access to difficult-to-collect and expensive data, hedge fund trades convey information. This may not have been as obvious to a trader in 1990. Furthermore, limits on attention may cause the public release of information to have a gradual impact on asset prices (e.g., Choi and Sias (2012)).

8 Conclusion

This paper provides novel evidence that hedge funds predict and drive the movement of asset prices towards fundamental value.

I apply the intuition of market microstructure models to the quarterly investment behavior of hedge funds. In particular, following the intuition of the Kyle model that price impact is a function of volume, I construct a measure of information that scales hedge fund purchases by total volume. If large trades relative to volume cause price impact, then fund managers should only be willing to consume a large share of volume when their information is especially compelling. Indeed, I find that the volume consumed by quarterly hedge fund trades strongly predicts future stock returns. Volume consumed also predicts earnings returns and comoves with insider trades. These results suggest that hedge funds are informed.

I confirm further predictions of the Kyle model to make the case that the price impact of hedge fund trades incorporates information into asset prices. Hedge fund trades appear to impound earnings information into prices *prior* to the information's public release: the impact of a given positive earnings surprise is reduced by 28% for stocks in the top quintile of volume consumed. I also show that volume consumed is positively associated with contemporaneous returns and predicts future trading, and that these price movements do not revert over multi-year horizons.

Large mutual fund trades are significantly less informative. However, the volume consumed by the trades of subsets of plausibly skilled mutual funds does predict future returns.

I estimate the quantitative price impact function using its reduced form and the full structural model. I find that consuming 1% of quarterly volume generates 0.3%-0.5% of permanent price impact. The model generates simulated moments of trading and returns that are reasonably close to the corresponding empirical moments.

My results highlight that one must examine asset prices before a fund's first trade to properly account for the information that a fund incorporates into prices. Due to price impact, prices move away from funds as they build large positions. The post-purchase prices of investment holdings and fund-level returns do not fully account for this effect. A fund with poor returns based on these metrics could still be identifying a substantial amount of information and helping to incorporate that information into prices.

I provide trade-level support for decreasing returns to scale in active management. I do so by showing that a portion of the intraday price impact documented in the microstructure literature aggregates at quarterly time-scales. Quarterly price movements are more relevant to many of the economic decisions of firm managers.

References

[1] Agarwal, Vikas, Vyacheslav Fos, and Wei Jiang, 2013, Inferring reporting-related biases in hedge fund databases from hedge fund equity holdings, Management Science 59, 1271-1289.

[2] Agarwal, Vikas, Wei Jiang, Yuehua Tang, and Baozhong Yang, 2013, Uncovering hedge fund skill from the portfolio holdings they hide, Journal of Finance 68, 739-783.

[3] Anderson, Anne-Marie and Edward Dyl, 2005, Market structure and trading volume, Journal of Financial Research 28, 115-131.

[4] Baker, Malcolm, Lubomir Litov, Jessica Wachter and Jeffrey Wurgler, 2010, Can mutual fund managers pick stocks? Evidence from their trades prior to earnings announcements, Journal of Financial and Quantitative Analysis 45, 1111-1131.

[5] Back, Kerry, Henry Cao, Gregory Willard, 2000, Imperfect competition among informed traders, Journal of Finance 55, 2117-2155.

[6] Berk, Jonathan and Richard Green, 2004, Mutual fund flows and performance in rational markets, Journal of Political Economy 112, 1269-1295.

[7] Berk, Jonathan and Jules van Binsbergen, 2015, Measuring skill in the mutual fund industry, Journal of Financial Economics, forthcoming.

[8] Boulatov, Alex, Terrence Hendershott and Dmitry Livdan, 2013, Informed trading and portfolio returns, Review of Economic Studies 80, 35-72.

[9] Brennan, Michael and Avanidhar Subrahmanyam, 1996, Market microstructure and asset pricing: On the compensation for illiquidity in stock returns, Journal of Financial Economics 41, 441-464.

[10] Cao, Charles, Yong Chen, William Goetzmann, and Bing Liang, 2015, The role of hedge funds in the security price formation process, Working paper.

[11] Campbell, John, Tarun Ramadorai, and Allie Schwartz, 2009, Caught on tape: Institutional trading, stock returns, and earnings announcements, Journal of Financial Economics 92, 66-91.

[12] Chan, Louis and Josef Lakonishok, 1993, Institutional trades and intraday stock price behavior, Journal of Financial Economics 33, 173-199.

[13] Chan, Louis and Josef Lakonishok, 1995, The behavior of stock prices around institutional trades, Journal of Finance 50, 1147-1174

[14] Chen, Joseph, Harrison Hong, Ming Huang, and Jeffrey Kubik, 2004, Does fund size erode mutual fund performance? The role of liquidity and organization, American Economic Review 94, 1276-1302.

[15] Choi, Nicole and Richard Sias, 2012, Why does financial strength forecast stock returns? Evidence from subsequent demand by institutional investors, Review of Financial Studies 25, 1550-1587.

[16] Cohen, Lauren, Andrea Frazzini, and Christopher Malloy, 2008, The Small World of Investing: Board Connections and Mutual Fund Returns, Journal of Political Economy 116, 951–979.

[17] Cohen, Lauren, Dong Lou, and Christopher Malloy, 2014, Playing favorites: How firms prevent the revelation of bad news, Working paper.

[18] Cohen, Randolph, Christopher Polk, and Bernhard Silli, 2010, Best ideas, Working paper.

[19] Collin-Dufresne, Pierre and Vyacheslav Fos, 2015, Do prices reveal the presence of informed trading?, Journal of Finance, forthcoming.

[20] Coval, Joshua and Erik Stafford, 2007, Asset fire sales (and purchases) in equity markets, Journal of Financial Economics 86, 479-512.

[21] Cremers, Martijn, and Antti Petajisto, 2009, How active is your fund manager? A new measure that predicts performance, Review of Financial Studies 22, 3329-3365.

[22] Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers, 1997, Measuring mutual fund performance with characteristic-based benchmarks, Journal of Finance 52, 1035–1058

[23] Del Guercio, Diane and Jonathan Reuter, 2014, Mutual fund performance and the incentive to generate alpha, Journal of Finance 69, 1673-1704.

[24] Foster, F. Douglas and S. Viswanathan, 1996, Strategic trading when agents forecast the forecasts of others, Journal of Finance 51, 1437-1478.

[25] Frazzini, Andrea and Owen Lamont, 2008, Dumb money: Mutual fund flows and the cross-section of stock returns, Journal of Financial Economics 88, 299-322.

[26] Griffin, John and Jin Xu, 2009, How smart are the smart guys? A unique view from hedge fund stock holdings, Review of Financial Studies 22, 2532-2570.

[27] Grinblatt, Mark, Sheridan Titman, and Russ Wermers, 1995, Momentum investment strategies, portfolio performance, and herding: A study of mutual fund behavior, American Economic Review 85, 1088-1105.

[28] Hendershott, Terrence, Dmitry Livdan, and Norman Schurhoff, 2015, Are institutions informed about news?, Journal of Financial Economics, forthcoming.

[29] Holden, Craig and Avanidhar Subrahmanyam, 1992, Long-lived private information and imperfect competition, Journal of Finance 47, 247-270.

[30] Hong, Harrison, Weikai Li, Sophie Ni, Jose Scheinkman, and Philip Yan, 2015, Days to cover and stock returns, Working paper.

[31] Huberman, Gur and Werner Stanzl, 2004, Price manipulation and quasi-arbitrage, Econometrica 72, 1247-1275.

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

[33] Jurek, Jakub and Erik Stafford, 2015, The Cost of Capital for Alternative Investments, Journal of Finance 70, 2185-2226.

[34] Kacperczyk, Marcin, Clemens Sialm, and Lu Zheng, 2005, On the industry concentration of actively managed equity mutual funds, Journal of Finance 60, 1983–2011.

[35] Kacperczyk, Marcin, Clemens Sialm, and Lu Zheng, 2008, Unobserved actions of mutual funds, Review of Financial Studies 21, 2379-2416.

[36] Keim, Donald and Ananth Madhavan, 1995, Anatomy of the trading process: Empirical evidence on the behavior of institutional traders, Journal of Financial Economics 37, 371-398.

[37] Keim, Donald and Ananth Madhavan, 1996, The upstairs market for large-block transactions: Analysis and measurement of price effects, Review of Financial Studies 9, 1-36.

[38] Khan, Mozaffar, Leonid Kogan, George Serafeim, 2012, Mutual fund trading pressure: Firm-level stock price impact and timing of SEOs, Journal of Finance 67, 1371-1395.
[39] Kruttli, Mathias, Andrew Patton and Tarun Ramadorai, 2014, The impact of hedge funds on asset markets, Working paper.
[40] Koudijs, Peter, 2014, Those who know most: Insider trading in 18th c. Amsterdam, Working paper.
[41] Kumar, Alok and Charles Lee, 2006, Retail investor sentiment and return comovements, Journal of Finance 61, 2451-2486.
[42] Kyle, Albert, 1985, Continuous auctions and insider trading, Econometrica 53, 1315-1335.
[43] Lakonishok, Josef, Andrei Shleifer, Richard Thaler, and Robert Vishny, 1991, Window dressing by pension fund managers, American Economic Review Papers and Proceedings 81, 227-231.
[44] Lou, Dong, 2012, A flow-based explanation for return predictability, Review of Financial Studies 25, 3457-3489.
[45] Di Mascio, Rick, Anton Lines and Narayan Naik, 2015, Alpha decay, Working paper.
[46] Pastor, Lubos and Robert Stambaugh, 2012, On the size of the active management industry, Journal of Political Economy 120, 740-781.
[47] Pastor, Lubos, Robert Stambaugh, and Lucian Taylor, 2015a, Scale and skill in active management, Journal of Financial Economics 116, 23-45.
[48] Pastor, Lubos, Robert Stambaugh, and Lucian Taylor, 2015b, Do funds make more when they trade more?, Working paper.
[49] Patton, Andrew, Tarun Ramadorai, and Michael Streatfield, 2015, Change you can believe in? Hedge fund data revisions, Journal of Finance, 70, 963-999.
[50] Pollet, Joshua and Mungo Wilson, 2008, How does size affect mutual fund behavior?, Journal of Finance 63, 2941-2969.
[51] Rhinesmith, Jonathan, 2014, Doubling Down, Working paper.
[52] Shleifer, Andrei and Robert Vishny, 1997, The limits of arbitrage, Journal of Finance 52, 35-55.
[53] Wurgler, Jeffrey and Ekaterina Zhuravskaya, 2002, Does arbitrage flatten demand curves for stocks?, Journal of Business 75, 583-608.

Figures

Figure 1: Timeline of Kyle model hypotheses for informed volume

This figure displays the chronological incidence of my hypotheses regarding informed volume. These hypotheses are derived from a two-period Kyle model of price impact (Section 2).

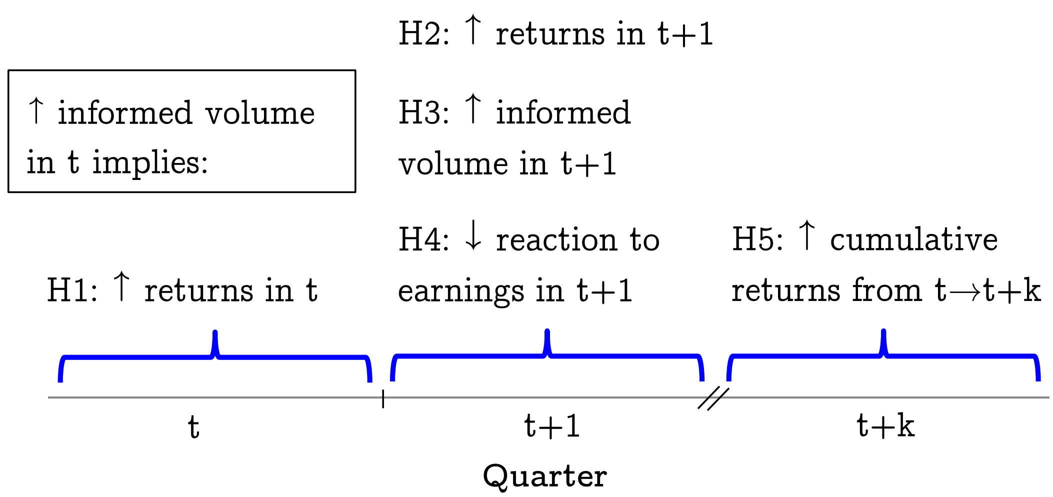

Figure 2: Volume consumed, cumulative returns

This figure displays the cumulative buy and hold performance of portfolios that go long stocks in the top decile of hedge fund volume consumed (aggregation method 1, see Section 3.2) and short stocks in the lowest decile. Calculations are based on 13F filings from 12/31/1989-9/30/2012 and use calendar-time portfolios (see Section 4.5). Panel A displays future market-adjusted performance, while Panel B displays future characteristic-adjusted performance. Panel C displays market-adjusted performance and also includes contemporaneous performance. Panel D displays characteristic-adjusted performance and also includes contemporaneous performance.

Panel A: Future market-adjusted returns

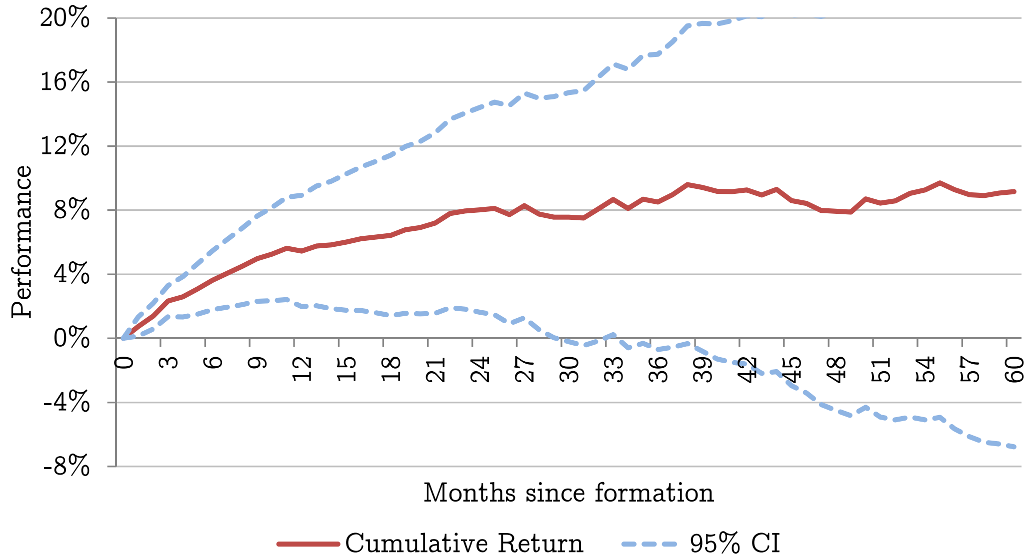

Panel B: Future characteristic-adjusted returns

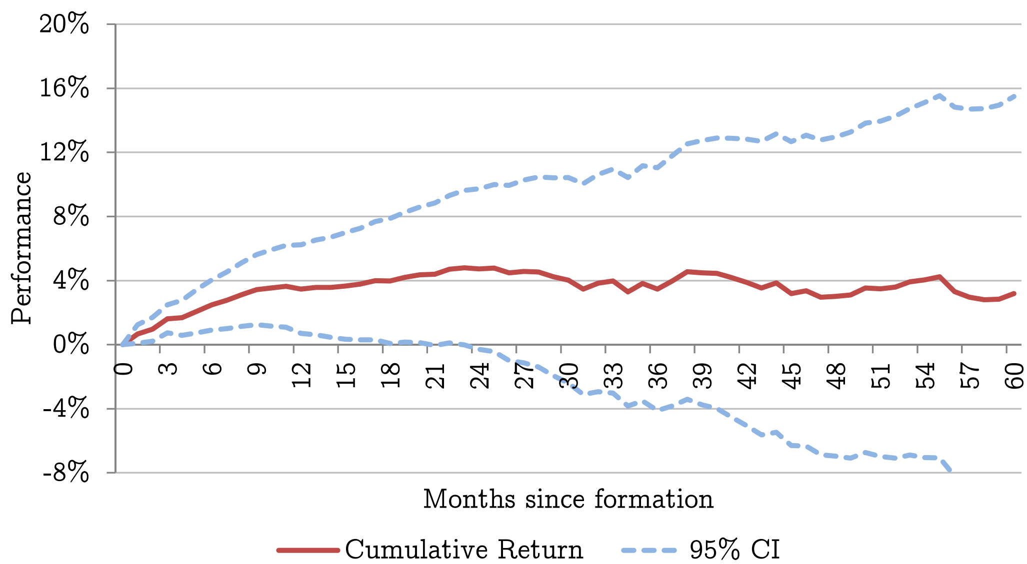

Panel C: Contemporaneous and future market-adjusted returns

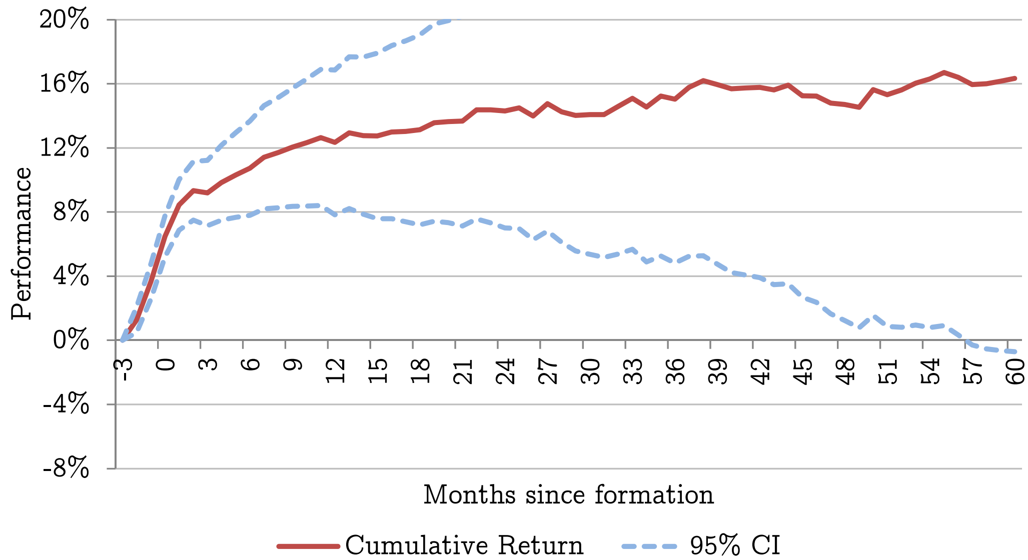

Panel D: Contemporaneous and future characteristic-adjusted returns

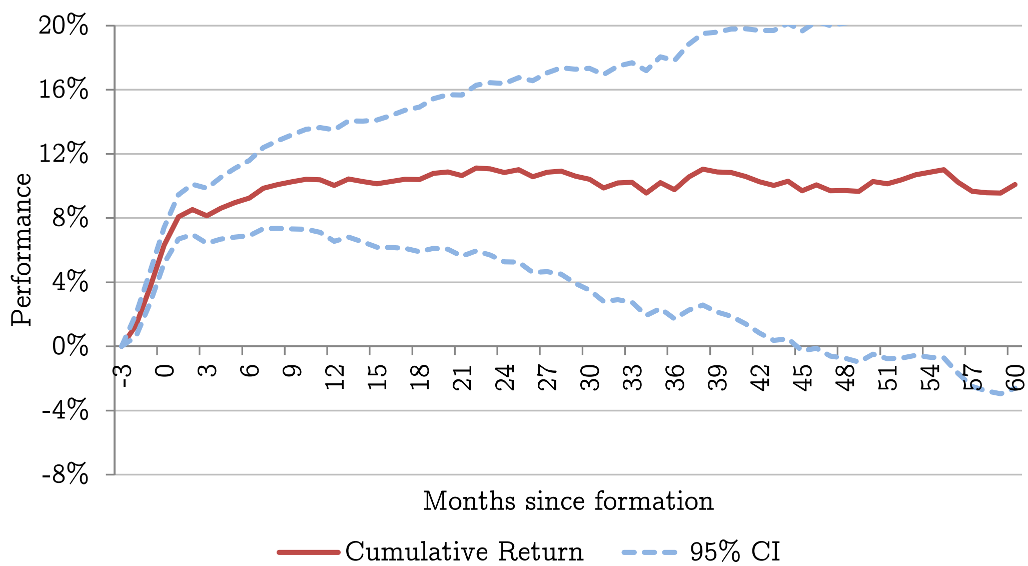

Figure 3: Mutual fund volume consumed, cumulative returns

This figure displays the cumulative buy and hold performance of portfolios that go long stocks in the top decile of mutual fund volume consumed (aggregation method 1) and short stocks in the lowest decile. Calculations are based on mutual funds' reported holdings from 12/31/1989-9/30/2012 and use calendar-time portfolios (see Section 4.5). Panel A includes only future returns, while Panel B also includes contemporaneous returns.

Panel A: Future market-adjusted returns

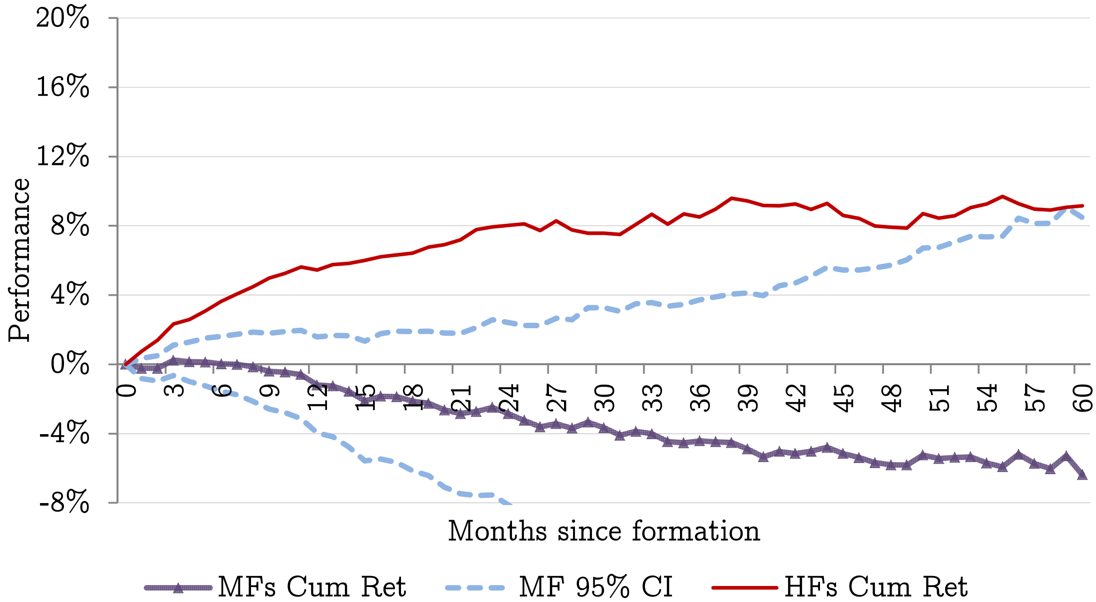

Panel B: Contemporaneous and future market-adjusted returns

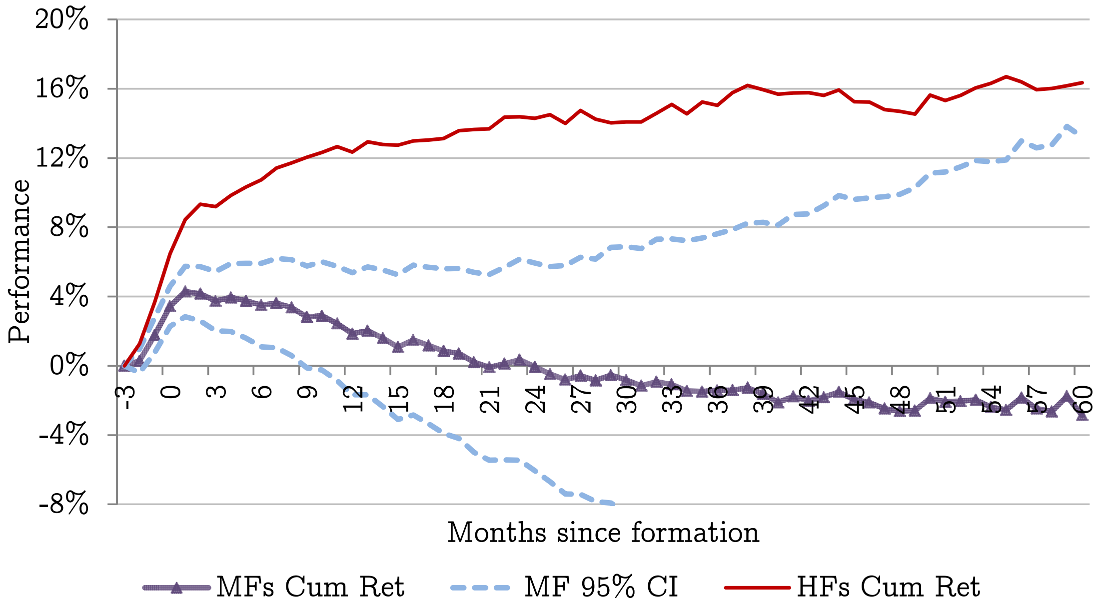

Tables

Table 1: Summary statistics

This table displays summary statistics of the hedge fund sample and of hedge fund volume consumed portfolios by decile. Calculations are based on 13F filings from 12/31/1989-9/30/2012. Statistics are calculated as the time-series average across 13F filings. Panel A presents summary statistics of the full sample. At each date, averages are calculated as the equal-weighted average across managers. For each manager, characteristic quintile averages are calculated using portfolio weights. A value of 5 represents a higher measure of the underlying statistic, i.e., the largest market cap quintile, the highest book-to-market quintile, or the highest trailing 12-month performance (excluding the most recent month) quintile. Panel B presents information on volume consumed portfolios by decile (aggregation method 1, see Section 3.2). The figure at each date is calculated as the equal-weighted average or total sum across all positions in the underlying portfolios described in the text. In panel B, volume consumed has been winsorized at the 1%/99% levels, and stocks below the 20th percentile of NYSE market cap have been removed.

Panel A: Hedge fund universe summary statistics

	Mean	Median	10th pctl	90th pctl	Standard Deviation
Hedge funds per quarter	308	276	84	574	187.4
Total positions per quarter	14,681	15,590	5,201	23,555	7,162
Total long U.S. equity assets per quarter ($ BB)	$231.2	$178.6	$40.8	$524.9	$183.6
Avg position size quintile	4.0	4.0	3.9	4.2	0.1
Avg position book quintile	2.8	2.8	2.6	2.9	0.1
Avg position momentum quintile	3.2	3.2	2.9	3.4	0.2

Panel B: Volume consumed summary statistics

Decile of volume consumed (t)	Average of volume consumed (t) (% quarterly vol)	Number of stocks	Number of hedge funds per stock	Total value of trades ($ BB)	Median stock mkt cap ($ BB)	Median stock volume ($ BB)
1	0.04%	142	2.1	$0.19	$2.06	$0.88
2	0.15%	143	3.8	$1.03	$2.28	$1.14
3	0.29%	143	4.7	$2.14	$2.20	$1.21
4	0.48%	143	4.7	$2.82	$1.99	$1.12
5	0.74%	143	4.9	$3.74	$1.82	$1.04
6	1.12%	143	4.8	$4.42	$1.60	$0.90
7	1.68%	143	4.8	$5.38	$1.41	$0.78
8	2.62%	143	4.8	$6.52	$1.23	$0.67
9	4.50%	143	4.9	$8.43	$1.08	$0.57
10	9.33%	142	5.3	$12.31	$0.86	$0.32

Table 2: Contemporaneous performance

This table displays in columns 1-5 estimated coefficients using contemporaneous market-adjusted and characteristic-adjusted monthly returns during quarter t as dependent variables. The explanatory variable is hedge fund volume consumed in quarter t by aggregation methods 1, 2, and 3. VCQ is the volume consumed quintile (1-5 for stocks with hedge fund trades, and 0 for stocks with no hedge fund trades) for stock s during quarter t. For comparison, in columns 6 and 7, I also display a regression to predict quarter t+1 returns using a stock's quintile of quarter t characteristic-adjusted returns (method †; 1-5, with a higher number corresponding to a higher return during quarter t) as an explanatory variable, and a regression of quarter t returns on a stock's quintile of quarter t opening value of trades, $valoftrade_{s,t}^{open} = shares\,traded_{s,t} * P_{s,t-1}$ (method ‡; 1-5, with a higher number corresponding to a higher opening value of trades), respectively. All variables are winsorized at the 1%/99% levels. Calculations are based on 13F filings from 12/31/1989-9/30/2012. T-statistics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively.

Column:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Agg. method:	(1)	(1)	(1)	(2)	(3)	(†)	(‡)
Dependent variable	Char.-adj ret (t)	Mkt.-adj ret (t)	Char.-adj ret (t)	Char.-adj ret (t)	Char.-adj ret (t)	Char.-adj ret (t+1)	Char.-adj ret (t)
VCQ (t)	0.28% [8.90]**	0.32% [6.74]**	0.33% [10.24]**	0.18% [7.25]**	0.29% [10.07]**		
Char.-adj return quintile (t)						0.01% [0.14]	
Opening value of trade quintile (t)							-0.13% [-4.51]**
Constant	-0.38% [-4.12]**	-0.25% [-1.08]	-0.23% [-2.32]**	0.19% [1.99]**	-0.54% [-5.04]**	0.12% [0.47]	0.98% [7.04]**
Fama-MacBeth	Y	Y	Y	Y	Y	Y	Y
Only volume consumed≠0	Y	Y	-	Y	Y	-	Y
Observations	170,384	195,610	408,924	676,121	217,912	408,924	170,384
R-squared	0.007	0.007	0.006	0.004	0.005	0.008	0.002

Table 3: Future performance

This table displays the future market-adjusted and characteristic-adjusted monthly performance of stocks during quarter t+1 based on hedge fund volume consumed in quarter t by aggregation methods 1, 2, and 3. VCQ is the volume consumed quintile (1-5 for stocks with hedge fund trades, and 0 for stocks with no hedge fund trades) for stock *s* during quarter *t*. Calculations are based on 13F filings from 12/31/1989-9/30/2012. T-statistics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively. Panel A displays results using Fama-MacBeth regressions with monthly returns during quarter t+1 as the dependent variable. In panel A, all variables are winsorized at the 1%/99% levels. Panel B displays monthly returns during quarter t+1 for calendar-time decile portfolios. In panel B, positions are weighted equally and stocks below the 20th percentile of NYSE market cap have been removed.

Panel A. Future performance – regressions

Column:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Agg. method:	(1)	(1)	(1)	(2)	(2)	(3)	(3)
Dependent variable	Char.-adj ret (t+1)	Mkt.-adj ret (t+1)	Char.-adj ret (t+1)	Char.-adj ret (t+1)	Mkt.-adj ret (t+1)	Char.-adj ret (t+1)	Mkt.-adj ret (t+1)
VCQ (t)	0.13%	0.17%	0.13%	0.09%	0.14%	0.11%	0.13%
	[4.73]**	[4.95]**	[5.15]**	[4.18]**	[3.19]**	[6.47]**	[7.71]**
Constant	-0.23%	-0.24%	-0.24%	-0.09%	-0.16%	-0.25%	-0.21%
	[-2.18]**	[-1.40]	[-2.58]**	[-0.97]	[-1.41]	[-2.97]**	[-1.17]
Fama-MacBeth	Y	Y	Y	Y	Y	Y	Y
Only volume consumed≠0	Y	Y	-	Y	Y	Y	Y
Observations	170,384	195,610	408,924	676,121	746,160	217,912	252,503
R-squared	0.003	0.004	0.003	0.002	0.006	0.001	0.001

Panel B. Future performance – portfolios

Column:	(1)	(2)	(3)	(4)	(5)	(6)
Agg. method:	(1)	(1)	(2)	(2)	(3)	(3)
Decile of volume consumed (t)	Char.-adj ret (t+1)	Mkt.-adj ret (t+1)	Char.-adj ret (t+1)	Mkt.-adj ret (t+1)	Char.-adj ret (t+1)	Mkt.-adj ret (t+1)
1	-0.08%	-0.04%	0.07%	0.08%	0.08%	0.22%
	[-1.04]	[-0.28]	[1.17]	[0.96]	[1.08]	[1.51]
2	-0.07%	-0.02%	0.05%	0.04%	-0.07%	0.01%
	[-0.86]	[-0.14]	[0.88]	[0.54]	[-1.23]	[0.10]
3	-0.04%	0.01%	-0.01%	-0.02%	-0.10%	-0.09%
	[-0.59]	[0.04]	[-0.11]	[-0.17]	[-1.50]	[-0.68]
4	-0.11%	-0.03%	0.03%	0.04%	-0.02%	0.05%
	[-1.59]	[-0.25]	[0.39]	[0.42]	[-0.23]	[0.39]
5	0.12%	0.16%	0.09%	0.13%	-0.08%	-0.03%
	[1.60]	[1.22]	[1.33]	[1.18]	[-1.28]	[-0.22]
6	0.21%	0.29%	0.11%	0.15%	0.02%	0.04%
	[2.84]**	[1.94]*	[1.83]*	[1.39]	[0.30]	[0.30]
7	0.25%	0.38%	0.16%	0.27%	0.09%	0.14%
	[3.14]**	[2.54]**	[2.71]**	[2.37]**	[1.25]	[1.09]
8	0.26%	0.35%	0.26%	0.38%	0.22%	0.32%
	[3.38]**	[2.45]**	[4.87]**	[3.15]**	[3.25]**	[2.33]**
9	0.37%	0.54%	0.28%	0.43%	0.30%	0.47%
	[4.85]**	[3.40]**	[5.27]**	[3.30]**	[4.32]**	[3.15]**
10	0.47%	0.70%	0.39%	0.59%	0.49%	0.68%
	[5.76]**	[4.84]**	[6.50]**	[4.37]**	[6.99]**	[4.70]**
L/S (10-1)	0.55%	0.74%	0.32%	0.51%	0.42%	0.46%
	[4.56]**	[5.36]**	[3.48]**	[2.84]**	[4.59]**	[5.10]**

Table 4: Future trading

This table displays estimated coefficients using measures of hedge fund volume consumed in quarter t+1 as dependent variables. Volume consumed is expressed as a percent of lagged quarterly volume. VCQ is the volume consumed quintile (1-5 for stocks with hedge fund trades, and 0 for stocks with no hedge fund trades) for stock s during quarter t or $t+1$, as specified. In each regression, volume consumed in quarter t+1 (dependent variable) is calculated using the same aggregation method used to calculate volume consumed during quarter t (explanatory variable). Calculations are based on 13F filings from 12/31/1989-9/30/2012. T-statistics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively. Volume consumed has been winsorized at the 1%/99% levels.

Column:	(1)	(2)	(3)	(4)	(5)
Agg. method:	(1)	(1)	(1)	(2)	(3)
Dependent variable	Volume consumed % (t+1)	VCQ (t+1)	Volume consumed % (t+1)	Volume consumed % (t+1)	Volume consumed % (t+1)
VCQ (t)	0.75%	0.26	0.57%	0.28%	0.41%
	[30.21]**	[31.43]**	[36.21]**	[22.83]**	[19.04]**
Constant	-0.01%	1.45	0.65%	-0.41%	-1.06%
	[-0.29]	[91.75]**	[19.79]**	[-23.06]**	[-10.36]**
Fama-MacBeth	Y	Y	Y	Y	Y
Only volume consumed (t) $\neq 0$	Y	Y	-	Y	Y
Observations	195,610	195,610	511,692	746,160	252,503
R-squared	0.060	0.049	0.069	0.041	0.006

Table 5: SUE and earnings returns

This table displays estimated coefficients involving earnings announcement returns and earnings surprises. VCQ is the volume consumed quintile (aggregation method 1; 1-5 for stocks with hedge fund purchases, and 0 for stocks with no hedge fund purchases) for stock s during quarter t. The characteristic-adjusted earnings return measures the return of stock s during the three trading-day window centered around its first earnings announcement during a quarter. SUE is the standardized earnings surprise for stock s in quarter $t+1$, defined as $\frac{earnings_{s,t+1} - median\ analyst\ forecast_{s,t}}{P_{s,t}}$, normalized to have a cross-sectional standard deviation of one each quarter. The characteristic-adjusted non-earnings return measures the daily return during a given quarter for stock s across all days except for the three trading-day earnings window, multiplied by three for comparability. $ME_{s,t}$, $V^{-1}_{s,t-1}$, $IOR_{s,t}$, and $BEME_{s,t}$ are the log of market cap, the log of the inverse of dollar volume, the level of institutional ownership, and the log of the book-to-market ratio of stock s at the end of quarter t (t-1 for volume), respectively. $var1 * var2$ is an interaction of $var1$ and $var2$. All variables are winsorized at the 1%/99% levels. Calculations are based on 13F filings from 12/31/1989-9/30/2012. T-statistics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively. Panel A examines how SUE and volume consumed predict earnings returns, and the relationship between SUE and volume consumed. Panel B focuses on the interaction of VCQ and SUE for observations with positive SUE. Panel C repeats the analysis of panel B for observations with negative SUE.

Panel A: How SUE and volume consumed predict earnings returns

Column:	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable	Char.-adj earnings ret (t+1)	Char.-adj non-earnings ret (t+1)	Char.-adj earnings ret (t+1)	Char.-adj earnings ret (t+1)	Char.-adj earnings ret (t+1)	SUE (t+1)
VCQ (t)	0.07%	0.02%	0.06%		0.03%	0.017
	[4.49]**	[5.98]**	[3.24]**		[1.56]	[10.38]**
SUE (t+1)				2.25%	2.24%	
				[10.78]**	[10.88]**	
$ME_{s,t}$	0.00%	0.01%	-0.63%	-0.32%	-0.32%	-0.164
	[0.04]	[0.36]	[-7.91]**	[-4.11]**	[-4.03]**	[-19.91]**
$V^{-1}_{s,t-1}$	0.12%	0.02%	-0.25%	-0.07%	-0.06%	-0.100
	[2.09]**	[0.93]	[-3.65]**	[-0.92]	[-0.83]	[-16.78]**
$IOR_{s,t}$	0.51%	-0.07%	0.01%	0.78%	0.73%	-0.379
	[3.27]**	[-2.03]**	[0.04]	[4.40]**	[4.08]**	[-19.63]**
$BEME_{s,t}$	-0.10%	-0.01%	0.06%	-0.23%	-0.23%	0.164
	[-2.52]**	[-1.83]*	[1.38]	[-5.27]**	[-5.18]**	[23.27]**
Fama-MacBeth	Y	Y	Y	Y	Y	Y
Only SUE≥0	-	-	Y	Y	Y	Y
Observations	120,749	120,749	80,362	80,362	80,362	80,362
R-squared	0.009	0.030	0.022	0.040	0.042	0.147

Panel B. Interaction of volume consumed quintile and positive SUE

Column:	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	Char.-adj earnings ret (t+1)	Char.-adj earnings ret (t+1)	Char.-adj earnings ret (t+1)	Char.-adj earnings ret (t+1)	Char.-adj earnings ret (t+1)	Char.-adj earnings ret (t+1)
SUE (t+1) *VCQ (t)	-0.14%	-0.20%	-0.29%	-0.21%	-0.13%	-0.13%
	[-2.16]**	[-3.46]**	[-2.41]**	[-2.11]**	[-3.63]**	[-3.54]**
SUE (t+1)	2.52%	4.52%	4.08%	1.78%	-0.60%	-2.57%
	[8.59]**	[0.99]	[0.77]	[0.44]	[-0.61]	[-2.50]**
VCQ (t)	0.07%	0.09%	0.11%	0.12%	0.08%	0.12%
	[2.55]**	[3.26]**	[3.18]**	[2.99]**	[4.23]**	[5.66]**
$ME_{s,t}$	-0.31%	-0.45%	-0.45%	-0.42%	-0.55%	-1.00%
	[-4.04]**	[-4.91]**	[-4.12]**	[-3.84]**	[-5.43]**	[-5.43]**
$V_{s,t-1}^{-1}$	-0.06%	-0.24%	-0.23%	-0.26%	-0.29%	-0.47%
	[-0.82]	[-2.81]**	[-2.41]**	[-2.51]**	[-3.44]**	[-4.14]**
$IOR_{s,t}$	0.74%	-0.72%	-0.60%	-0.92%	0.17%	-0.09%
	[4.13]**	[-2.12]**	[-2.11]**	[-2.71]**	[0.72]	[-0.23]
$BEME_{s,t}$	-0.24%	-0.08%	-0.03%	0.05%	0.02%	0.27%
	[-5.28]**	[-1.11]	[-0.46]	[0.65]	[0.26]	[3.02]**
Fama-MacBeth	Y	Y	Y	Y	-	-
Controls interacted w/SUE	-	Y	Y	Y	Y	Y
No bottom quintile mkt cap	-	-	Y	-	-	-
Only volume consumed≠0	-	-	-	Y	-	-
Std err clustered by firm	-	-	-	-	Y	-
Std err clustered by time	-	-	-	-	Y	Y
Firm fixed effects	-	-	-	-	-	Y
Observations	80,362	80,362	69,774	59,151	80,362	80,362
R-squared	0.044	0.057	0.056	0.064	0.028	0.082

Panel C. Interaction of volume consumed quintile and negative SUE

Column:	(1)	(2)	(3)	(4)
Dependent variable:	Char.-adj earnings ret (t+1)	Char.-adj earnings ret (t+1)	Char.-adj earnings ret (t+1)	Char.-adj earnings ret (t+1)
SUE (t+1) *VCQ (t)	-0.03%	-0.02%	-0.04%	-0.01%
	[-1.66]*	[-0.88]	[-1.22]	[-0.38]
SUE (t+1)	0.45%	1.02%	2.20%	1.05%
	[7.98]**	[1.21]	[1.76]*	[0.86]
VCQ (t)	0.13%	0.14%	0.14%	0.22%
	[5.01]**	[5.06]**	[5.12]**	[6.08]**
$ME_{s,t}$	0.35%	0.32%	0.31%	0.23%
	[3.89]**	[3.44]**	[3.14]**	[2.26]**
$V^{-1}_{s,t-1}$	0.18%	0.18%	0.14%	0.03%
	[2.34]**	[2.21]**	[1.52]	[0.29]
$IOR_{s,t}$	-1.10%	-1.17%	-1.11%	-0.99%
	[-4.20]**	[-4.04]**	[-3.96]**	[-2.92]**
$BEME_{s,t}$	0.50%	0.53%	0.54%	0.52%
	[6.55]**	[6.88]**	[6.53]**	[5.86]**
Fama-MacBeth	Y	Y	Y	Y
Controls interacted w/SUE	-	Y	Y	Y
No bottom quintile mkt cap	-	-	Y	-
Only volume consumed≠0	-	-	-	Y
Observations	40,387	40,387	33,595	27,265
R-squared	0.040	0.058	0.064	0.073

Table 6: Mutual fund trades and performance

This table displays the monthly performance of stocks based on mutual fund volume consumed in quarter t. VCQ is the volume consumed quintile (aggregation method 1; 1-5 for stocks with mutual fund purchases, and 0 for stocks with no mutual fund purchases) for stock s during quarter t. All variables are winsorized at the 1%/99% levels. Calculations are based on mutual funds' reported holdings from 12/31/1989-9/30/2012 (except for active share results, which end at 12/31/2009). T-statistics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively. Panel A analyzes the contemporaneous (quarter t) and future (quarter $t+1$) monthly performance of mutual fund trades. Panel B analyzes the future monthly performance of the trades of subsets of mutual funds: the top (bottom) quintile of return gap in column 1 (column 2) and funds with above (below) median active share in column 3 (column 4).

Panel A: Mutual fund contemporaneous and future performance

Column:	(1)	(2)	(3)	(4)
Dependent variable	Char.-adj ret (t)	Mkt.-adj ret (t)	Char.-adj ret (t+1)	Mkt.-adj ret (t+1)
VCQ (t)	0.17%	0.17%	0.01%	-0.01%
	[5.12]**	[4.36]**	[0.49]	[-0.54]
Constant	0.03%	0.23%	0.08%	0.26%
	[0.33]	[1.43]	[0.90]	[1.56]
Fama-MacBeth	Y	Y	Y	Y
Only volume consumed≠0	Y	Y	Y	Y
Observations	111,664	125,129	111,664	125,129
R-squared	0.005	0.005	0.002	0.002

Panel B: Mutual fund subsets – future performance

Column:	(1)	(2)	(3)	(4)
Subset:	Return gap, top quintile	Return gap, bottom quintile	Active share > median	Active share < median
Dependent variable	Char.-adj ret (t+1)	Char.-adj ret (t+1)	Char.-adj ret (t+1)	Char.-adj ret (t+1)
VCQ (t)	0.05%	0.02%	0.04%	0.00%
	[1.77]*	[0.61]	[1.66]*	[-0.12]
Constant	-0.02%	0.04%	-0.02%	0.06%
	[-0.19]	[0.38]	[-0.16]	[0.56]
Fama-MacBeth	Y	Y	Y	Y
Only volume consumed≠0	Y	Y	Y	Y
Observations	79,702	76,481	110,022	60,751
R-squared	0.003	0.003	0.002	0.002

Table 7: Quantifying the price impact function

This table displays structural and reduced-form estimates – based on hedge fund trading – of parameters of the Kyle model as well as moments of informed trading and returns. The structural model is estimated via maximum likelihood. The model splits all observations into two-quarter intervals, ending at March 31 and September 30 (baseline) or June 30 and Dec 31 ("oth date"). The model is estimated on 10 or 100 portfolios formed after sorting by hedge fund volume consumed (aggregation method 1). I include estimates based on an informed trader with constant absolute risk aversion ("RA") and estimates that feature both risk aversion and public new information shocks ("RA+NI"). The reduced form is estimated by Fama-MacBeth regressions, without splitting observations into intervals. The average percent of information incorporated into prices during the first and second quarters of trading implied by the model – $\frac{average(p_1)}{average(\epsilon)}$ and $\frac{average(p_2-p_1)}{average(\epsilon)}$, respectively, within the censored simulated data – is also displayed. x_1 signifies hedge fund volume consumed. r_1 signifies *quarterly* characteristic-adjusted returns. Standard errors are displayed in parentheses, based on time-clustered bootstraps (structural model) or Fama-MacBeth regressions (reduced form). Calculations are based on 13F filings from 12/31/1989-9/30/2012. Panel A displays estimates of the coefficient of price impact, λ_1. It also displays the first and second moments of informed trading and returns of model-simulated and empirical observations with greater than 0.1% of volume consumed, the censoring cutoff. Panel B displays other parameters from the structural model.

Panel A. Kyle's λ and moments

	10 portf.	10 portf., oth date	100 portf.	100 portf., RA	100 portf., RA+NI	10 portf.	100 portf.	Full cross section
	Structural model					Reduced form		
λ_1	0.32	0.33	0.53	0.43	0.31	0.29	0.30	0.31
	(0.03)	(0.03)	(0.04)	(0.04)	(0.04)	(0.03)	(0.03)	(0.03)
	Model-implied moments					Empirical moments		
x_1	4.6%	4.0%	5.5%	6.7%	6.1%	3.8%	3.8%	3.9%
	(0.3%)	(0.2%)	(0.3%)	(0.3%)	(0.5%)			
σ_{x1}	3.4%	3.0%	4.1%	5.0%	4.6%	5.5%	6.1%	6.6%
	(0.2%)	(0.1%)	(0.2%)	(0.3%)	(0.4%)			
r_1	1.5%	1.3%	2.9%	2.9%	1.8%	2.1%	2.1%	1.8%
	(0.1%)	(0.1%)	(0.2%)	(0.2%)	(0.5%)			
σ_{r1}	2.6%	2.4%	4.9%	4.9%	5.7%	3.5%	6.5%	23.9%
	(0.2%)	(0.2%)	(0.3%)	(0.3%)	(0.4%)			
Observations	506	506	4,646	4,646	4,646	1,012	9,292	400,413
Time periods	46	46	46	46	46	92	92	92

Panel B. Structural parameters

	10 portf.	10 portf., oth date	100 portf.	100 portf., RA	100 portf., RA+NI
	Structural model				
% of information incorporated into prices during quarter 1	37.9%	36.7%	39.6%	38.7%	36.4%
% of information incorporated into prices during quarter 2	30.4%	30.7%	29.7%	30.4%	30.8%
λ_2	0.26	0.28	0.42	0.34	0.25
	(0.03)	(0.02)	(0.03)	(0.03)	(0.03)
β_1	1.18	1.11	0.76	0.91	1.26
	(0.13)	(0.09)	(0.06)	(0.08)	(0.12)
β_2	1.92	1.83	1.20	1.39	2.04
	(0.20)	(0.15)	(0.10)	(0.12)	(0.21)
σ_ϵ	4.8%	4.5%	9.1%	9.2%	6.0%
	(0.4%)	(0.4%)	(0.6%)	(0.6%)	(0.4%)
σ_u	7.2%	6.5%	8.4%	10.4%	9.7%
	(0.4%)	(0.3%)	(0.5%)	(0.6%)	(0.8%)
π	46.4%	46.0%	46.4%	46.4%	46.4%
	(0.5%)	(0.6%)	(0.5%)	(0.5%)	(0.5%)
σ_η	0.0%	0.0%	0.0%	0.0%	0.0%
	(0.0%)	(0.0%)	(0.0%)	(0.0%)	(0.0%)
Observations	506	506	4,646	4,646	4,646
Time periods	46	46	46	46	46

Table 8: Insider trades

This table displays the results of regressions involving the trades of firm insiders and monthly characteristic-adjusted returns during quarter $t+1$. VCQ is the volume consumed quintile (aggregation method 1; 1-5 for stocks with hedge fund purchases, and 0 for stocks with no hedge fund purchases) for stock s during quarter t. "Insider purchase?" is an indicator variable equal to 1 if firm insiders were net purchasers of stock s during quarter t, and 0 otherwise. $ME_{s,t}$, $V_{s,t-1}^{-1}$, $IOR_{s,t}$, and $BEME_{s,t}$ are the log of market cap, the log of the inverse of dollar volume, the level of institutional ownership, and the log of the book-to-market ratio of stock s at the end of quarter t (t-1 for volume), respectively. All variables are winsorized at the 1%/99% levels. Calculations are based on 13F filings from 12/31/1989-9/30/2012. T-statistics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively.

Column	(1)	(2)	(3)	(4)
Dependent variable	Char.-adj ret (t+1)	Char.-adj ret (t+1)	Insider purchase? (t)	Char.-adj ret (t+1)
VCQ (t)	0.14%		0.002	0.14%
	[7.60]**		[2.43]**	[7.64]**
Insider purchase? (t)		0.59%		0.58%
		[9.07]**		[8.85]**
$ME_{s,t}$	0.08%	0.08%	0.020	0.07%
	[0.59]	[0.61]	[9.52]**	[0.51]
$V_{s,t-1}^{-1}$	-0.14%	-0.11%	-0.019	-0.13%
	[-1.20]	[-1.01]	[-13.63]**	[-1.12]
$IOR_{s,t}$	0.66%	0.91%	-0.099	0.72%
	[2.52]**	[3.49]**	[-15.41]**	[2.74]**
$BEME_{s,t}$	0.06%	0.05%	0.022	0.05%
	[1.27]	[0.93]	[12.01]**	[1.01]
Fama-MacBeth	Y	Y	Y	Y
Only volume consumed≠0	-	-	-	-
Observations	328,778	328,778	328,778	328,778
R-squared	0.027	0.027	0.029	0.028

Table 9: Idiosyncratic risk-weights and 13F dates

This table displays the characteristic-adjusted monthly performance during quarter t+1 of portfolios either formed based on hedge fund volume consumed and idiosyncratic risk weights during quarter t or based on hedge fund volume consumed during quarter t but with performance split by the 13F filing date during quarter t+1. Calculations are based on 13F filings from 12/31/1989-9/30/2012. Positions are weighted equally. T-statistics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively. Stocks below the 20th percentile of NYSE market cap have been removed. Panel A compares volume consumed and best ideas. Positive volume consumed (aggregation method 2) positions are sorted into quintiles, and then bucketed into three groups: positions with no volume consumed or in the bottom quintile; positions in the middle three quintiles; and positions in the top quintile. Positions are independently sorted by their intra-manager best ideas ranking (relative to other stocks s in fund f's portfolio at quarter t). The proportion of total positions within each bin is displayed in italics. Panel B displays performance during quarter t+1 of volume consumed (aggregation method 1) portfolios split by the 13F filing date the following quarter (45 days after the previous quarter end) into three time periods: before the 13F window (monthly performance), the three-trading-day window centered around the 13F date (absolute performance), and after the 13F window (monthly performance).

Panel A: Volume consumed and best ideas

Char.-adj ret (t+1) / [t-stat] / *proportion of total positions*)

Best ideas position rank (t; 1 = highest best ideas)

		11+	4-10	1-3
Volume consumed (t)	None or bottom quintile	0.06%	0.06%	0.13%
		[2.41]**	[0.86]	[1.15]
		49.0%	*5.9%*	*2.4%*
	Middle quintiles	0.13%	0.06%	0.07%
		[3.00]**	[0.59]	[0.48]
		26.4%	*4.5%*	*1.7%*
	Top quintile	0.30%	0.42%	0.35%
		[5.33]**	[5.28]**	[2.92]**
		6.5%	*2.3%*	*1.2%*

Panel B: 13F filing dates

Column:	(1)	(2)	(3)	(4)
	Before 13F	During 13F	After 13F	After 13F-Before 13F
Decile of volume consumed (t)	Char.-adj ret (t+1)	Char.-adj ret (t+1)	Char.-adj ret (t+1)	Char.-adj ret (t+1)
1	-0.04%	0.02%	-0.14%	-0.10%
	[-0.39]	[0.79]	[-1.31]	[-0.69]
10	0.50%	0.03%	0.44%	-0.06%
	[4.23]**	[0.93]	[3.63]**	[-0.36]
L/S (10-1)	0.54%	0.01%	0.58%	0.06%
	[3.25]**	[0.24]	[3.17]**	[0.18]

Appendix

A Proofs

A.1 Additional model equations

The following equilibrium equations constrain the parameters of the model:

$$\beta_1 = \frac{1}{2\lambda_1}\left(\frac{\lambda_2 - \frac{1}{2}\lambda_1\pi}{\lambda_2 - \frac{1}{4}\lambda_1\pi}\right) \tag{6}$$

$$\beta_2 = \frac{1}{2\lambda_2} \tag{7}$$

$$\lambda_1 = \frac{\beta_1\phi\sigma_\epsilon^2}{\beta_1^2\phi^2(\sigma_\epsilon^2+\sigma_\eta^2)+\sigma_u^2} \tag{8}$$

$$\lambda_2 = \frac{\beta_2\phi(1-\lambda_1\beta_1\phi)\sigma_\epsilon^2}{\beta_2^2\phi^2(1-\lambda_1\beta_1\phi)(\sigma_\epsilon^2+\sigma_\eta^2)+\sigma_u^2} \tag{9}$$

$$\phi = \frac{\sigma_\epsilon^2}{\sigma_\epsilon^2+\sigma_\eta^2} \tag{10}$$

A.2 Model solution

To solve the model, I conjecture equations (1)-(4) and (6)-(10). I then verify the equilibrium.

1. The informed agent solves for her optimal trading pattern given the price impact function set by the market maker.
   (a) In $t = 2$, solve $max_{x_2} E[x_2(\epsilon - p_2)\,|\,i]$. This gives (2), with $\beta_2 = \frac{1}{2\lambda_1}$ as in (7). The second order condition implies $\lambda_2 > 0$.
   (b) In $t = 1$, solve $max_{x_1} E[x_1(\epsilon - p_1)\,|\,i] + \pi E[x_2(\epsilon - p_2)\,|\,i]$. This gives (1), with β_1 as in (6). The second order condition implies $\lambda_1 > 0$.
2. The market maker attempts to infer ϵ from order flow, given the informed trader's (optimized) behavior. That is, the market maker sets $p_1 = E[\epsilon|\,y_1]$ and $p_2 = E[\epsilon|\,y_1,\,y_2]$.
   (a) Solve for p_1 using Bayesian updating. The result is (3), with λ_1 as in (8).
   (b) Solve for p_2 also using Bayesian updating and assuming that y_1 and y_2 are independent. The result is (4), with λ_2 as in (9). Note further that (9) can be written as $\lambda_2 = \frac{\beta_2\phi(1-\lambda_1\beta_1\phi)\sigma_\epsilon^2}{\beta_2^2\phi^2(1-\lambda_1\beta_1\phi)(\sigma_\epsilon^2+\sigma_\eta^2)+\sigma_u^2} = \sqrt{\frac{\phi(1-\lambda_1\beta_1\phi)\sigma_\epsilon^2}{4\sigma_u^2}}$.
   (c) Confirm that $cov(y_1,\,y_2) = 0$, which is used in the calculation of the price rule in (2.b). This covariance follows by plugging in and noting that $(1-\lambda_1\beta_1)\beta_1^2\phi\sigma_\epsilon^2 = \lambda_1\beta_1\sigma_u^2$. Because y_1 and y_2 are jointly normally distributed, that is sufficient for independence.

A.3 Hypothesis (1) : $cov(p_1,\, x_1) > 0$

This result follows immediately because $p_1 = \lambda_1(x_1 + u_1)$, with u_1 independent of all other first period random variables.

A.4 Hypothesis (2) : $cov(p_2 - p_1,\, x_1) > 0$ and $cov(\epsilon - p_1,\, x_1) > 0$

For the first covariance, rewrite $p_2 - p_1 = \lambda_2(x_2 + u_2) = \lambda_2\beta_2(\frac{1}{\beta_1}x_1 - \lambda_1 x_1 - \lambda_1 u_1) + \lambda_2 u_2$. Discarding terms with u_t, since those are constant or independent of x_1, leaves $\frac{1}{2\beta_1}(1 - \beta_1\lambda_1)x_1$. $\beta_1\lambda_1 \leq \frac{1}{2}$ because $\beta_1\lambda_1 = \frac{\lambda_2 - \frac{1}{2}\pi\lambda_1}{2\lambda_2 - \frac{1}{2}\pi\lambda_1}$ and $\lambda_1 > 0$ and $\beta_1 > 0$. The latter follows because $\lambda_2 > 0$ and $\lambda_1 > 0$ from the informed trader's second order conditions, and $\beta_1 > 0$ to satisfy (7). Thus $cov(p_2 - p_1,\, x_1) = \frac{1}{2\beta_1}(1 - \beta_1\lambda_1)cov(x_1, x_1) > 0$.

For the second covariance, use $x_1 = \beta_1\phi(\epsilon + \eta)$ and $\epsilon - p_1 = \epsilon - \lambda_1 x_1 - \lambda_1 u_1 = \epsilon - \lambda_1\beta_1\phi(\eta + \epsilon) - \lambda_1 u_1$. Then $cov(\epsilon - \lambda_1\beta_1\phi(\eta + \epsilon) - \lambda_1 u_1,\, x_1) = \beta_1\phi\sigma_\epsilon^2 - \lambda_1\beta_1^2\phi^2\sigma_\epsilon^2 - \lambda_1\beta_1^2\phi^2\sigma_\eta^2$, since u_1 is independent of ϵ and η. Plug in for λ_1 to get $cov(\epsilon - p_1,\, x_1) = \beta_1\phi\sigma_\epsilon^2(1 - \frac{\beta_1^2\phi^2(\sigma_\epsilon^2 + \sigma_\eta^2)}{\beta_1^2\phi^2(\sigma_\epsilon^2 + \sigma_\eta^2) + \sigma_u^2}) > 0$.

A.5 Hypothesis (3) : $cov(x_2,\, x_1) > 0$

Rewrite $cov(x_2, x_1) = cov(\frac{1}{\lambda_2}(p_2 - p_1) - u_2,\, x_1)$. $\lambda_2 > 0$, u_2 is independent of x_1, and $cov(p_2 - p_1,\, x_1) > 0$ as shown in Hypothesis (2). Therefore $cov(x_2,\, x_1) > 0$.

A.6 Hypothesis (4) : $\frac{\partial}{\partial\eta'}E(\epsilon - p_1|\eta = \eta', \epsilon > 0) < 0$.

Write out the expectation to obtain $E(\epsilon - p_1|\eta = \eta', \epsilon > 0) = \frac{2}{\pi}\sigma_\epsilon(1 - \lambda_1\beta_1\phi) - \lambda_1\beta_1\phi\eta'$. The derivative of this expression with respect to η' is $-\lambda_1\beta_1\phi < 0$.

Note that $\frac{\partial}{\partial\eta'}E(\epsilon - p_2|\eta = \eta', \epsilon > 0) < 0$, too. If one conceptualizes earnings being released after the second period, then this is the model counterpart. Write out this expectation to obtain $E(\epsilon - p_2|\eta = \eta', \epsilon > 0) = \frac{2}{\pi}\sigma_\epsilon(1 - \frac{1}{2}\phi - \frac{1}{2}\lambda_1\beta_1\phi) - \eta'(\frac{1}{2}\phi + \frac{1}{2}\lambda_1\beta_1\phi)$. The derivative of this expression with respect to η' is $-(\frac{1}{2}\phi + \frac{1}{2}\lambda_1\beta_1\phi) < 0$.

In mapping this equation to its empirical counterpart, I use the fact that as η increases, so does the probability of seeing a higher x_1. Therefore higher volume consumed (x_1) on average corresponds to higher η.

A.7 Hypothesis (5) : $cov(\bar{\epsilon}^K,\, x_1) > 0$, $K > 1$

This follows from the fact that future information draws have an expectation of zero ($E[\epsilon_{k'}|\, x_1] = 0$ for $k' > k$, where k is the current episode). The cumulative expected price movement from future episodes is thus zero. The price movement resulting from the current episode should persist, on average.

A.8 Additional Proof: $E[mispricing] = \phi i = constant * \frac{x_1}{\sigma_u}$

Rewrite (5) as $\lambda_1\beta_1 = \frac{\sqrt{a_1} - \pi\sigma_u\sqrt{\lambda_1\beta_1}}{2\sqrt{a_1} - \pi\sigma_u\sqrt{\lambda_1\beta_1}}$, with $a_1 = \sigma_u^2 + \beta_1^2\phi^2\sigma_\eta^2$. Plug in for λ_1 using (8) to eliminate all λ_1 terms. Obtain the equation $a_3 * (2\sqrt{a_1 a_2} - \pi\sigma_u\sqrt{a_3}) = a_2 * (\sqrt{a_1 a_2} - \pi\sigma_u\sqrt{a_3})$ with $a_2 = \sigma_u^2 + \beta_1^2\phi\sigma_\epsilon^2$ and $a_3 = \beta_1^2\phi\sigma_\epsilon^2$. Plugging in $\beta_1' = C\beta_1$ and $\sigma_u' = C\sigma_u$ also solves this equation, if β_1 and σ_u do. In other words, if you double expected noise trading in each period, you double β_1, i.e., how much the insider trades for a given amount of information.

Thus $x_1 = \beta_1 \phi i$ implies $E[mispricing] = \phi i = constant * \frac{x_1}{\sigma_u}$. Again, σ_u is a measure of the magnitude of expected noise trading, since the expectation of the absolute value of a normal random variable centered at zero is proportional to its standard deviation.

One can most easily see the mathematical intuition for this result from a one-period Kyle model. Using the notation of my model, start at the beginning of period 2. Since the insider knows the information will be revealed at the end of the period, the model proceeds as a one-period Kyle model from that point on. The "initial" price is p_1, for a mispricing of $\phi i - p_1$. Then $x_2 = \beta_2(\phi i - p_1) = constant * \sigma_u * (E[remaining\, mispricing])$, after simply plugging in for β_2 from (7) and (9). Thus $E[mispricing] = constant * \frac{x_2}{\sigma_u}$.

B Example and robustness to alternative explanations

B.1 Illustrative example

Baupost's purchase of IHOP during the year 2000 illustrates my basic approach. Entering the second quarter of the year 2000, the International House of Pancakes' stock (IHOP) had seen better days.[54] As of March 31, 2000, the stock had fallen 29% over the prior 12 months, underperforming the value-weighted market index by 55%.

IHOP drew the attention of Baupost, a well-known Boston-based hedge fund headed by Seth Klarman. Between April 1 and June 30, Baupost purchased 1.5 million shares of IHOP out of total volume of 2.7 million shares. Baupost was the buyer of *55% of the shares sold* during the quarter.[55] As Baupost purchased these shares, IHOP rose from \$14 to \$16.75, returning 19.6%.

By purchasing such a large fraction of volume, Baupost signaled that it had compelling information about IHOP. Indeed, IHOP rose from \$16.75 to \$19.13 the following quarter, a return of 14.2%. Baupost continued to purchase IHOP from July 1 to September 30, buying an incremental 0.4 million shares out of total volume of 2.2 million shares.

The stock did not give back these gains over subsequent quarters and years. Instead, IHOP continued to outperform. Over the five years subsequent to April 1, 2000 – the first day of the quarter in which Baupost began to purchase IHOP – the stock beat the value-weighted market index by 270%.

The precise figures in this example are extreme. Nevertheless, I find that hedge funds consume large fractions of total volume – into the double digits – with regularity. These

[54] As its name implies, the company runs a chain of restaurants specializing in breakfast foods. IHOP Corp. changed its name to DineEquity in 2007 following its acquisition of Applebee's.

[55] I construct volume consumed using lagged volume. Relative to lagged volume, $volconsumed_{IHOP,2000Q2} = 68\%$.

large trades predict future returns. Furthermore, I present evidence that the price impact of these trades incorporates information into prices as the trades occur.

B.2 Alternative explanations

Volume consumed predicts future performance beyond existing alternative explanations. Table B.1 shows future performance (quarter t+1) after controlling for each alternative. To present my results most succinctly, and since some alternative explanations may be non-monotonic, I focus on the performance of the top-decile (for one-dimensional sorts) or quintile (for double sorts) portfolios. In general, these tests either remove positions susceptible to the alternative explanation or employ 5x5 dependent double sorts. For double sorts, I sort first along the dimension of the alternative and then by volume consumed (aggregation method 1) to determine if the latter has incremental explanatory power.

The explanatory power of volume consumed remains after controlling for each of the following alternatives:

1. Downward sloping demand, in its simplest form, would suggest that hedge fund trades exert temporary price pressure. In that case, we would expect to see returns revert, as after flow-driven mutual fund trades. Figure 2 illustrates that there is no evidence of reversion.
2. Heterogeneous beliefs or segmented demand suggest that purchases move up the optimism / valuation of the marginal holder of an asset. In that case, returns revert *after* a hedge fund sells its position. Figure B.1 shows that hedge funds substantially reduce positions before the cumulative performance of high volume-consumed positions becomes insignificant.
3. To clearly differentiate my results from a simple method of examining portfolio weights, I double sort by positions' portfolio weights. This is similar to the approach in Section 7.2, but does not rely on a proxy for the cross-section of stock-level risk.
4. Activists use large investments relative to firms' market caps to exert corporate control to directly influence the value of firms. I control for this alternative with several tests. First, I remove all positions in which a fund owns more than 5% of the market cap of a company. Second, I remove the top decile of hedge funds by funds' average stake in the companies in their portfolios (some funds are activist funds). Third, I double sort by the stake of a stock's market cap held by hedge funds.
5. Concentrated funds may outperform by more (Kacperczyk, Sialm, and Zheng (2005)). I double sort by a fund's number of positions to control for this alternative.
6. Flow-driven investing suggests that managers who have done well recently ("hot hands") attract flows. Investing those flows drives up the prices of their holdings (Coval and

Stafford (2007) and Lou (2012)). Double sorting by fund-level past performance, past flows, or even *future* flows does not explain my results.

7. Volume on its own does not capture the predictive power of volume consumed. Double sorting by inverse volume as a percent of shares outstanding or by inverse dollar volume does not eliminate the significance of my results. It does reduce magnitudes, since it is positively correlated with volume consumed by construction. I also double sort by the stake of volume hedge funds have invested in a company: the sum of the shares held divided by share volume (rather than total shares outstanding).
8. Proxies for asymmetric information and liquidity do not explain my results. I test double sorts by the volatility of daily returns over the past three months, by bid/ask spreads, by Amihud ratios, and by PIN.[56]
9. There are theoretical reasons to consider price impact as a function of volume. One could instead consider price impact as a function of market cap, and divide shares traded by a stock's total shares outstanding – "market cap consumed" – rather than by its share volume. Market cap and volume are highly correlated, but volume consumed has explanatory power even after first sorting by market cap consumed. In unreported results, I find that the reverse is *not* true: after first sorting by volume consumed, "market cap consumed" does not have incremental explanatory power.
10. For robustness, I display future performance including all stocks (i.e., without eliminating the bottom quintile by market cap) and including only stocks with above-median NYSE market cap or dollar volume.
11. The outperformance of stocks with high hedge fund volume consumed is not limited to a narrow period of time. Figure B.2 displays the trailing 3-year average characteristic-adjusted performance of a long-short portfolio formed from the extreme deciles of volume consumed. The only period during which this long-short portfolio generated negative 3-year performance was during the early 1990s.

[56] PIN data is from is from Jefferson Duarte's webpage, and ends in 2004.

Figure B.1: Volume consumed, cumulative performance and sales over time

This figure displays the cumulative buy and hold performance of portfolios that go long stocks in the top decile of hedge fund volume consumed (aggregation method 1) and short stocks in the lowest decile. It also displays the average stake of market cap held by hedge funds for positions in the top-decile portfolio over time. Calculations are based on 13F filings from 12/31/1989-9/30/2012.

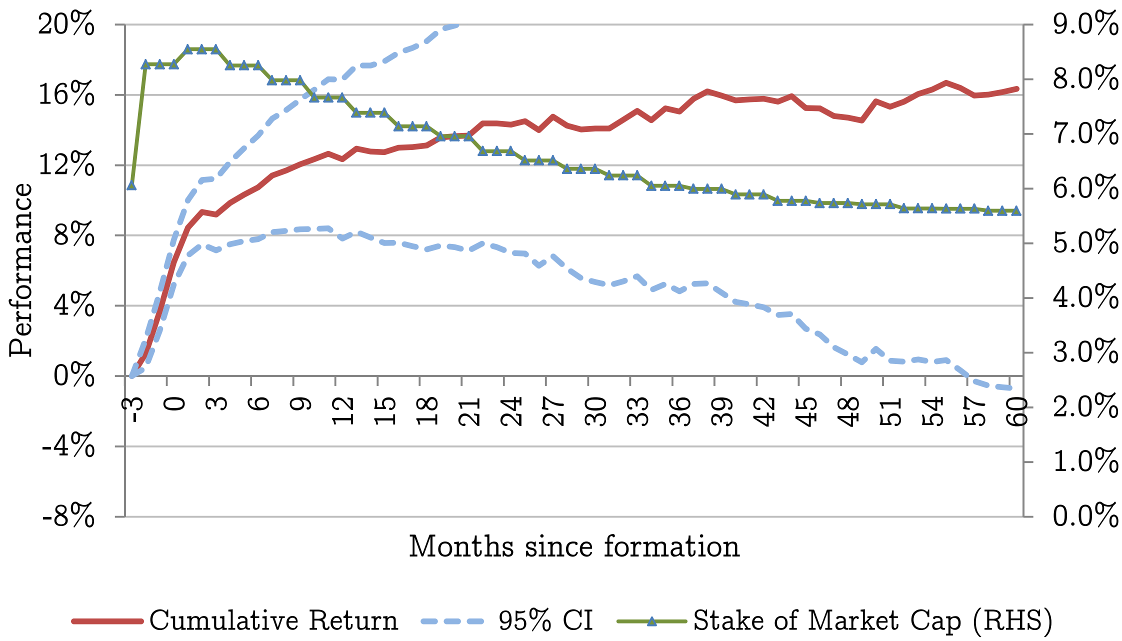

Figure B.2: Long-short portfolio, 3-year trailing characteristic-adjusted performance

This figure displays the trailing 3-year average monthly characteristic-adjusted performance of a portfolio that goes long stocks in the top decile of hedge fund volume consumed (aggregation method 1) and short stocks in the lowest decile. The portfolio is re-formed at the end of every quarter t based on volume consumed during quarter t, and is then held during quarter t+1. Calculations are based on 13F filings from 12/31/1987-9/30/2012 – as this figure emphasizes subsample performance, I include data from before 12/31/1989 to illustrate that the outperformance that I identify is not sensitive to the start date of my sample.

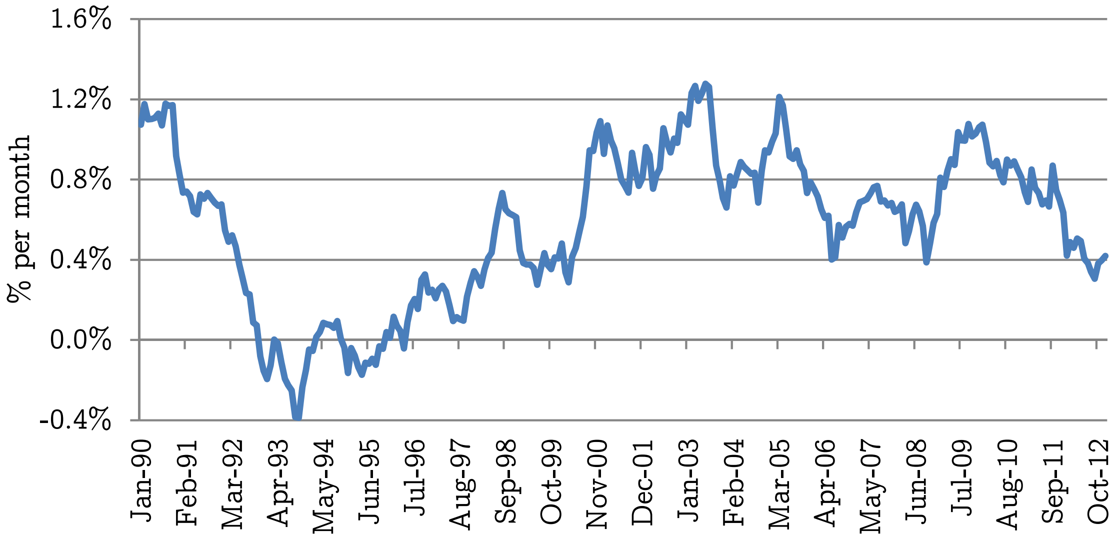

Table B.1: Alternative explanations

This table displays the characteristic-adjusted monthly performance during quarter t+1 of portfolios that compare hedge fund volume consumed in quarter t (aggregation method 1) to a variety of alternative empirical controls from quarter t (except for future flows, which are from quarter t+1) described in Appendix B.2. The estimates are of the monthly performance of the top quintile/decile portfolios, and of portfolios that go long the top quintile/decile portfolio and short the lowest quintile/decile portfolio, for the controls described in the text. Calculations are based on 13F filings from 12/31/1989-9/30/2012. Positions are weighted equally. T-statistics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively. Stocks below the 20th percentile of NYSE market cap have been removed.

Characteristic-adjusted performance (t+1)

Control	L/S portfolio	Top decile/ quintile	Control	L/S portfolio	Top decile/ quintile
			Volume		
Baseline result:	0.55%	0.47%	Double sort:	0.44%	0.30%
	[4.56]**	[5.76]**	stake of volume	[5.71]**	[5.32]**
Portfolio weights					
Double sort:	0.54%	0.47%	Double sort:	0.26%	0.33%
avg portfolio weight	[5.38]**	[7.25]**	inverse volume,	[4.11]**	[6.43]**
			% shares outstanding		
Activism					
No stake > 5% of	0.58%	0.51%	Double sort:	0.46%	0.38%
mkt cap	[4.70]**	[6.19]**	inverse volume, $ value	[5.84]**	[6.29]**
			Liquidity /		
No top decile			**asymmetric info**		
of managers by avg.	0.49%	0.46%	Double sort:	0.63%	0.46%
stake in company	[3.71]**	[5.31]**	volatility of daily returns	[4.33]**	[4.42]**
Double sort:	0.49%	0.35%	Double sort:	0.46%	0.42%
stake of mkt cap	[4.23]**	[6.11]**	bid/ask spreads	[3.40]**	[4.42]**
Concentration					
Double sort:	0.57%	0.52%	Double sort:	0.59%	0.43%
mgr avg no. of positions	[4.82]**	[6.21]**	Amihud ratio	[4.26]**	[4.16]**
Flows / hot hands					
Double sort:	0.49%	0.39%	Double sort:	0.39%	0.38%
trail 4 quarter perf	[4.49]**	[5.51]**	PIN	[3.87]**	[4.88]**
			Robustness		
Double sort:	0.49%	0.38%	No mkt cap filter	0.52%	0.43%
trail 4 quarter flows	[4.73]**	[5.53]**		[3.46]**	[4.06]**
Double sort:	0.46%	0.35%	Only mkt cap > median	0.41%	0.32%
future 1 quarter flows	[4.00]**	[4.78]**	NYSE	[3.45]**	[4.15]**
Normalize by mkt cap					
Double sort:	0.33%	0.26%	Only $ volume > median	0.74%	0.59%
Market cap consumed	[1.78]*	[4.10]**	NYSE	[5.86]**	[6.28]**