Essays on the Dynamic Strategies and Skill of Institutional Investors

Essays on the Dynamic Strategies and Skill of Institutional Investors

A dissertation presented

by

Jonathan Rhinesmith

to

The Department of Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

Harvard University

Cambridge, Massachusetts

May 2016

Dissertation Advisors:

Professor John Campbell

Professor Jeremy Stein

Author:

Jonathan Rhinesmith

Essays on the dynamic strategies and skill of institutional investors

Abstract

This dissertation studies the behavior of institutional investors, who control a large share of the world's investment capital, with the goal of shedding light on when and how those investors reveal information. Guided by economic intuition, I highlight instances in which the trades of fund managers are particularly informative. I focus on hedge funds and present evidence that in these instances funds' decisions predict future asset price movements.

These results demonstrate that fund managers possess valuable information. At the same time, my findings support a view of the world in which fund managers have more capital than what they allocate to opportunities with high expected returns. Hedge funds may be "smart" – they may be able to identify mispriced securities – while still delivering poor returns to their investors.

Chapter 1 presents evidence that price impact is an important consideration even at the quarterly time horizon of the trades I observe. If fund trades generate price impact, and if price impact is a function of volume, then funds should only be willing to trade a large share of volume when their information is compelling. Indeed, I find that hedge funds predict future stock returns when they purchase a large share of volume. I also provide evidence that the price impact of fund trades incorporates information into stock prices. If informative prices impact real economic decision making then these findings support the welfare relevance of the active management industry.

Chapter 2 shows that funds avoid adding to losing positions. When they do, however, they predict future stock-level outperformance. These results are consistent with a career risks mechanism, as adding to a losing position corresponds to reverse window dressing. They also suggest a position-level limits-to-arbitrage effect.

Chapter 3 demonstrates that hedge funds frequently buy back into stocks they have held in the past. This phenomenon occurs much more often than it would by chance. I use these findings to argue that fund managers develop company-specific expertise that persists over time. When funds establish expert positions after poor past stock-level performance, they predict future stock-level excess returns.

Contents

List of Tables

List of Figures

Acknowledgments

I am incredibly grateful to my advisors John Campbell, Lauren Cohen, Chris Malloy, and Jeremy Stein for their invaluable guidance. I also thank Malcolm Baker, Robin Greenwood, Sam Hanson, Owen Lamont, Jonathan Reuter, Andrei Shleifer, Erik Stafford, and Adi Sunderam for extremely helpful comments and suggestions. I thank seminar participants at Harvard, the Federal Reserve Board, the University of North Carolina at Chapel Hill, Rochester, Georgetown, London Business School, Northeastern, Notre Dame, Penn State, and Washington University in St. Louis for constructive feedback, and I thank my peers Divya Kirti, Filippo Mezzanotti, Vijay Narasiman, Pascal Noel, Thomas Powers, Yao Zeng, and John Zhou for many thoughtful discussions. I thank my family and close friends for supporting me throughout this process. Finally, I thank my former colleagues at Yale's Investments Office, including David Swensen and Dean Takahashi, for never letting me forget that the world is comprised of collections of individuals and institutions.

Introduction

In Chapter 1 of this dissertation, "Conviction and volume: Measuring the information content of hedge fund trading," 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.

In Chapter 2, "Doubling down," I demonstrate that when investment fund managers double down on positions that have run against them, those positions outperform. Specifically, I find that a portfolio formed of the U.S. equity positions that hedge fund managers add to after recent stock-level underperformance generates significant annualized risk-adjusted outperformance of between 5% and 15%. This finding is not the result of a simple reversal effect, of a fund's best ideas (large positions), or of the general informativeness of fund trades. My results are consistent with a career risks mechanism for this phenomenon. By adding to a losing position — the opposite of window dressing — managers are making their losses particularly salient. I demonstrate in a panel regression that investment managers avoid adding to losing positions. Furthermore, managers outperform by more when they double down after greater past losses in a position. These findings suggest a position-level limits-to-arbitrage effect. Even when an asset decreases in price for non-fundamental reasons, some of the investment managers with the most relevant knowledge of that asset may be particularly hesitant to add to their positions because they have already suffered losses in that asset.

In Chapter 3, "Stock experts," I show that many investment managers develop company-specific expertise about a subset of the firms within their potential investment universes. Within their long equity portfolios, hedge fund managers disproportionately reestablish large positions in stocks that they have held in the past. A portfolio of the expert positions purchased following stock-level underperformance generates abnormal returns of 5%-10% annualized, suggesting that fund managers are able to distinguish temporary stock price dislocations from fundamental company underperformance in those stocks. Comparable control portfolios of non-expert stocks fail to outperform. Furthermore, a higher level of available expert capital in a given stock predicts a substantial increase in return volatility. This finding could be driven by either information-based or non-fundamental factors. The existence of experts highlights the importance of the network of investment relationships in capital formation.

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

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 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

[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."

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

[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)).

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 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 1.1 reviews related literature. Section 1.2 develops a Kyle model of price impact and uses the model to generate testable hypotheses. Section 1.3 describes the data and constructs volume consumed. Section 1.4 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 1.5 shows that total mutual fund volume consumed is uninformative but that plausibly skilled subsets of funds reveal information through volume. Section 1.6 estimates the quantitative price impact function, both by employing the reduced form and by undertaking a structural estimation of the model. Section 1.7 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 1.8 concludes.

1.1. 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 1.2. 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

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

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).

1.2. 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.1.

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]

[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.

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. 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:

$$x_1 = \beta_1 \phi i \tag{1.1}$$

$$x_2 = \beta_2(\phi i - p_1) \tag{1.2}$$

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

$$p_1 = \lambda_1(x_1 + u_1) \tag{1.3}$$

$$p_2 = p_1 + \lambda_2(x_2 + u_2) \tag{1.4}$$

The equations determining the parameters λ_1, λ_2, β_1, and β_2 are detailed in Appendix A.1.1. 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 1.1 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

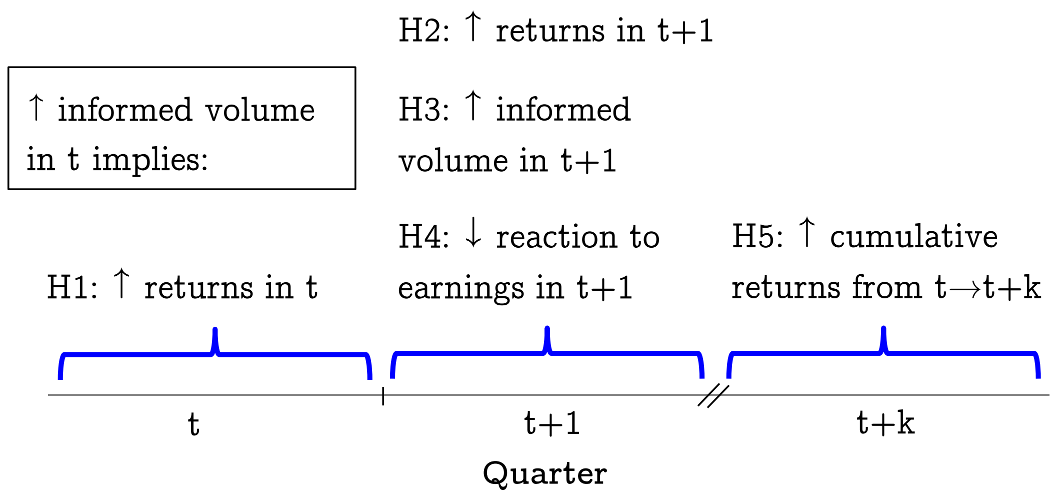

Figure 1.1. Timeline of Kyle model hypotheses for informed volume
This figure displays the chronological incidence of my hypotheses regarding informed volume. These hypotheses are derived from a two-period Kyle model of price impact (Section 1.2).

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 A.6.

1.3. Constructing volume consumed

1.3.1. 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

[6] More detailed requirements are provided at https://www.sec.gov/answers/form13f.htm. The full 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 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

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 1.7.3, 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

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

[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.

returns, trades, flows, and relative position sizes refer to the remaining 13F portfolios.

Table 1.1 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 with slight growth 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.

1.3.2. 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.1.8 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

[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.

Table 1.1. 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 1.3.2). 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

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 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 A.5.

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 in quarter t+1 – a standard test of skill and return predictability – and in Table 1.1 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

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

[16]In the other direction, in Appendix A.2.2 I demonstrate that limiting the sample to stocks above the NYSE median market cap or NYSE median dollar volume still produces highly significant results.

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 1.1 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 A.2.1).[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.

1.4. Volume consumed: price impact and information

1.4.1. 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 1.2 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

[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.

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 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 1.7.2).

Hedge funds are not naively purchasing stocks with high past risk-adjusted returns.

[19] Table A.3 illustrates this result using decile portfolios.

[20] See Section 1.6 for estimates of the price impact function and comparisons to existing estimates.

Table 1.2. 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, $valoftrade^{open}_{s,t} = shares\,traded_{s,t} * P_{s,t-1}$ (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

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 1.2 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 1.4.2, 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 $valoftrade^{open}_{s,t} = shares\,traded_{s,t} * P_{s,t-1}$, 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 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.

[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.

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 1.4.5, 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.

1.4.2. 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 1.3 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

Table 1.3. 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

Table 1.3: (continued)

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]**

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 1.7.2 and Appendix A.2.2, 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)

[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.

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).

1.4.3. 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 1.4 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

[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.

[27] Table A.4 illustrates this result using decile portfolios.

Table 1.4. 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

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 consumed in quarter t+1.

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

---

[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%.

[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.

1.4.4. 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.

1.4.4.1. 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

[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.

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:

$$\begin{aligned} earningsreturn_{s,t+1} &= \beta\,(VCQ_{s,t} * SUE_{s,t+1}) + \alpha_1\,VCQ_{s,t} + \alpha_2\,SUE_{s,t+1} \\ &\quad +\gamma\,controls_{s,t} + \nu_{s,t} \end{aligned} \tag{1.5}$$

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 A.4.1.

Table 1.5 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

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

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.

1.4.4.2. Reaction to positive SUE following informed purchases

Table 1.5 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 of the volume consumed quintile and SUE (equation (1.5)). 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%.

[32] Table A.5 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.

Table 1.5. 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

Table 1.5: (continued)

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

Table 1.5: (continued)
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]**
$V_{s,t-1}^{-1}$	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

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 (1.5) separately by volume consumed quintile, and find that the r-squared is highest for stocks in the top quintile of volume consumed.[33][34]

[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 *because* hedge funds trade more of those stocks. Nevertheless, first note that in Table 1.5 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

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.

1.4.4.3. Reaction to negative SUE following informed purchases

For completeness, Table 1.5 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

---

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.

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.

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

Figure 1.2 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

[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.

Panel A: Future market-adjusted returns

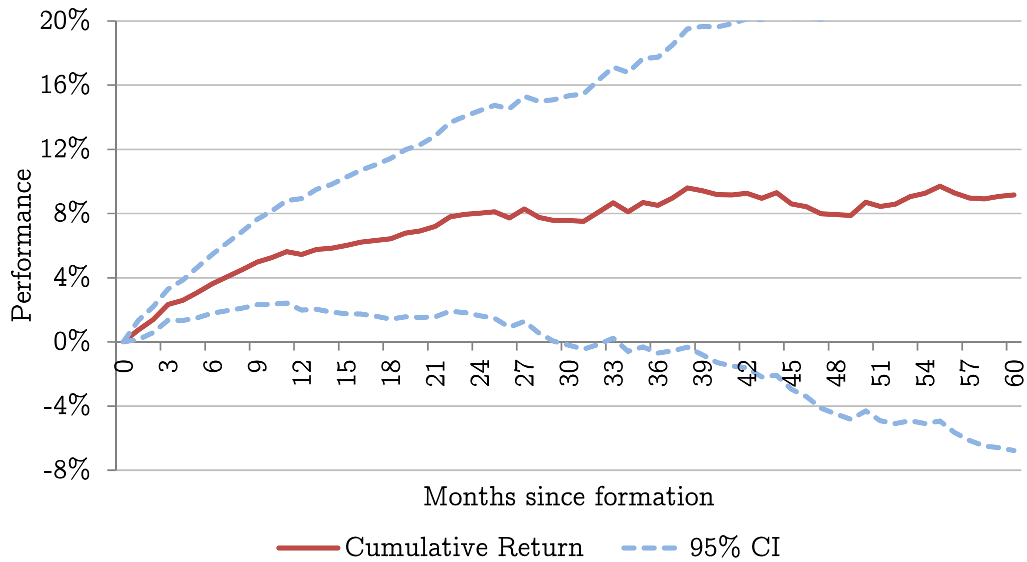

Panel B: Future characteristic-adjusted returns

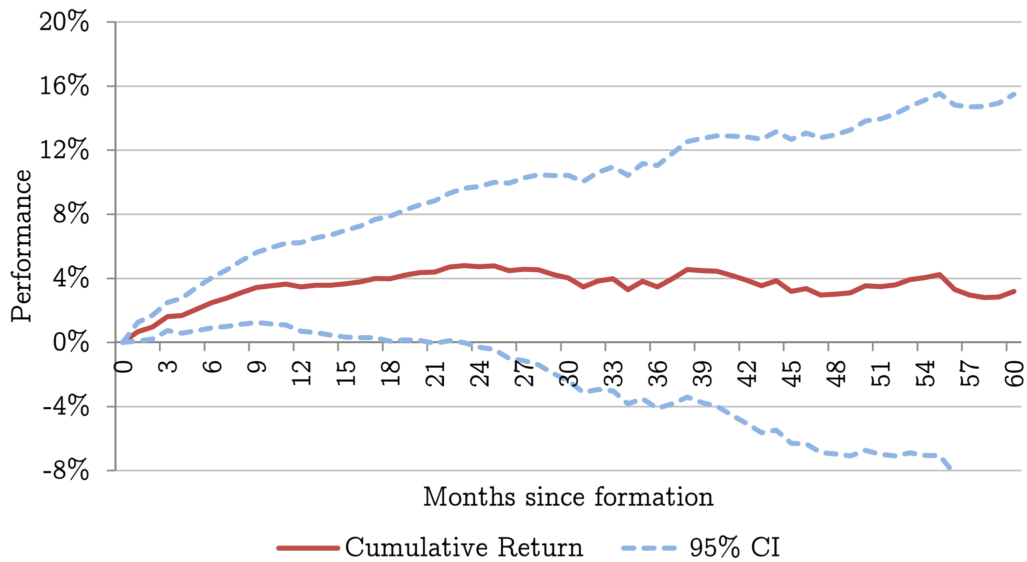

Figure 1.2. Volume consumed, cumulative returns

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 1.3.2) 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 1.4.5). 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 C: Contemporaneous and future market-adjusted returns

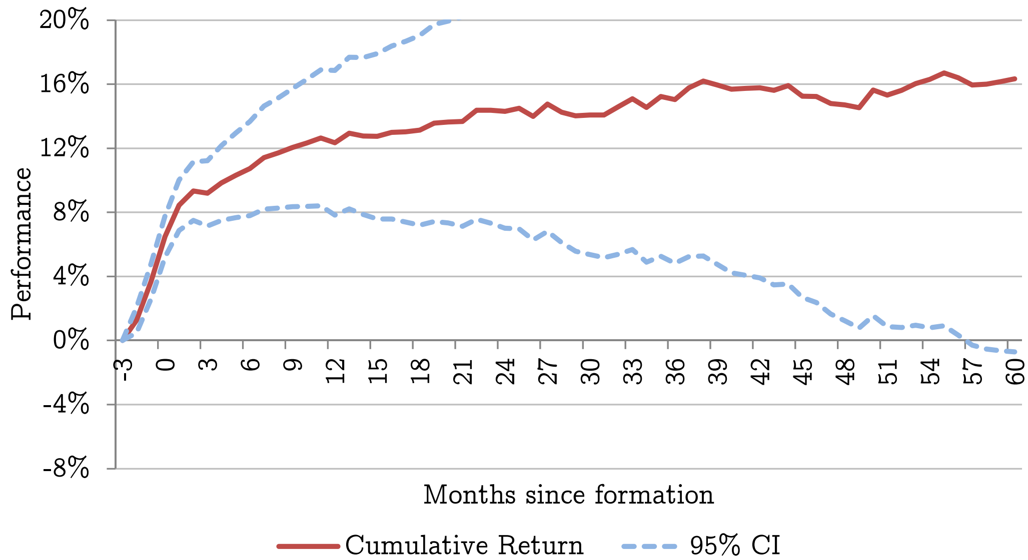

Panel D: Contemporaneous and future characteristic-adjusted returns

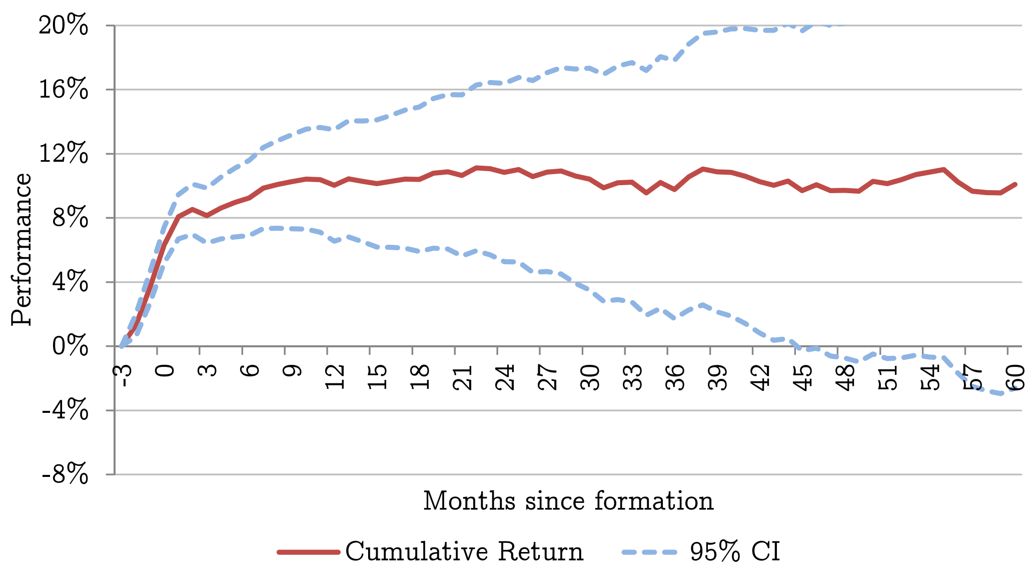

Figure 1.2: (continued)

about 16% (10%) on a market-adjusted (characteristic-adjusted) basis at a five-year horizon. Section 1.4.1 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 A.2.2). In contrast to these results, large mutual fund trades are associated with price movements that tend to revert – as in Section 1.5.1 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.

1.5. 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 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.

1.5.1. 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 A.4.2 for details). I construct volume consumed by aggregation method 1 (purchases aggregated at the stock level).[37]

Table 1.6 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-

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

[38]Table A.6 illustrates the results of this section using decile portfolios.

Table 1.6. 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% [5.12]**	0.17% [4.36]**	0.01% [0.49]	-0.01% [-0.54]
Constant	0.03% [0.33]	0.23% [1.43]	0.08% [0.90]	0.26% [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

Table 1.6: (continued)

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

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% (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 1.3, 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

Panel A: Future market-adjusted returns

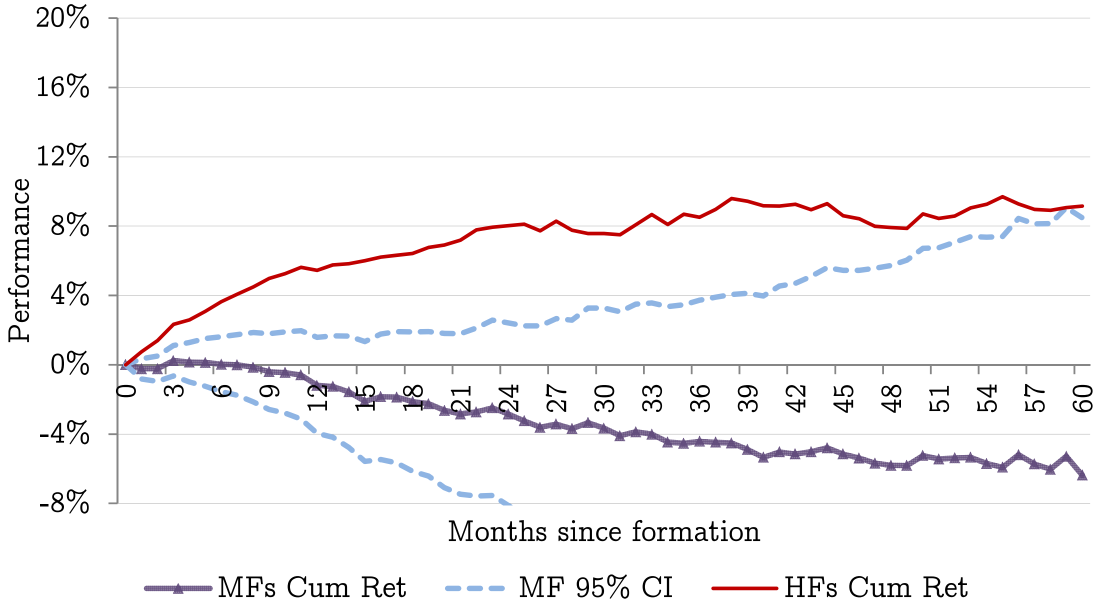

Panel B: Contemporaneous and future market-adjusted returns

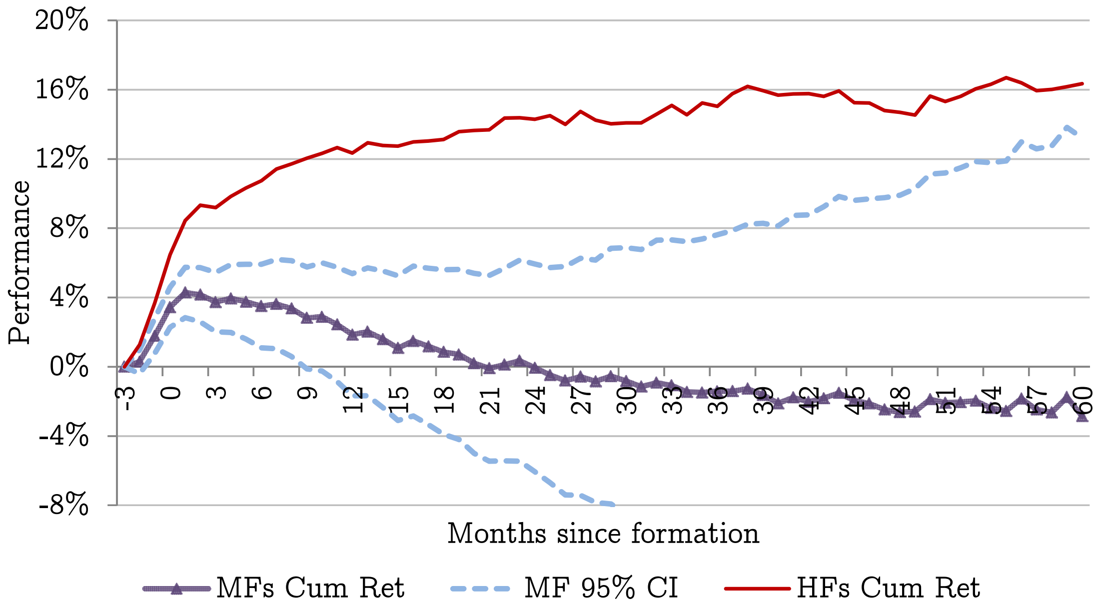

Figure 1.3. Mutual fund volume consumed, cumulative returns

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 1.4.5). Panel A includes only future returns, while Panel B also includes contemporaneous returns.

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.

1.5.2. 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 A.4.3 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 1.6 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 1.3).

1.6. Quantifying the price impact function

1.6.1. 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 1.2 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.

1.6.2. 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 (1.3) and (1.4). The model jointly fits empirical moments (1.1)-(1.4), subject to equilibrium constraints (A.1)-(A.5).

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 A.3.1. 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.

1.6.3. 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.

[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).

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 are due to trading. If higher average returns are associated with the same amount of trading,

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

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).

1.6.4. Estimates from hedge fund purchases

Table 1.7 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

[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 A.3.3, and are based on estimates from Frazzini, Israel, and Moskowitz (2012), Collin-Dufresne and Fos (2015), Brennan and Subrahmanyam (1996), and Investment Technology Group.

Table 1.7. 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

Table 1.7: (continued)

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

moments – equations (1.1) and (1.2) – 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 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 1.7, 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 1.6.3, the model interprets noise in returns as information. With less information but a similar

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

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 A.3.2) 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.

1.6.5. 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 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

[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.

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 1.4.4). 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.

[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.

1.7. Additional evidence

1.7.1. 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 1.8 examines how hedge fund and insider trades relate. I construct an indicator variable 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]

---

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

[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

Table 1.8. 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. $ME_{s,t}$, $V^{-1}_{s,t-1}$, $IOR_{s,t}$, 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. 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]
$V^{-1}_{s,t-1}$	-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

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.

1.7.2. 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 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

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.

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 1.9 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.

---

[50] Wurgler and Zhuravskaya (2002) motivate a similar approach. See Appendix A.5.2 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 A.7, 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.

Table 1.9. 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
	None or bottom quintile	0.06%	0.06%	0.13%
		[2.41]**	[0.86]	[1.15]
		49.0%	*5.9%*	*2.4%*
Volume consumed (t)	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%*

Table 1.9: (continued)

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]

In my hedge fund sample, volume consumed subsumes this measure of idiosyncratic risk.

1.7.3. 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 publicly observable. Table 1.9 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)).

1.8. 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.

2. Doubling down

> "If the security you are considering is truly a good investment, not a speculation, you would certainly want to own more at lower prices." Seth Klarman, Margin of Safety

Suppose an investment fund manager buys a stock for \$10 that she thinks is worth \$15. The stock proceeds to decline in value to \$7, while the market remains flat. As an econometrician, one cannot easily tell if the stock's fundamental value has dropped, or if the stock price movement was just noise, making the stock a better buy at \$7 than it was at \$10. If one believes the investment manager has skill, perhaps the investment manager can tell the difference. Yet even if the stock is a more attractive buy now, the fund manager may be hesitant to add to, or to "double down" on, her existing position. Her investors already know she has suffered substantial losses in the stock, and adding to the position will make those losses even more salient. The manager would effectively be employing reverse window dressing; instead of substituting out losing positions for winning stocks, she is making her losing positions even bigger. If such an effect were indeed at work, one would expect that fund managers would only "double down" in the most promising of situations, and that the corresponding positions would outperform.

In line with this reasoning, I find that in a sample of the long U.S. equity positions of hedge fund managers from January 1, 1990 through December 31, 2013, a portfolio formed of the positions that hedge fund managers add to following recent stock-level underperformance generates significant annualized risk-adjusted outperformance of between 5% and 15%. In turn, positions that managers double down on after greater position-level losses outperform by more than those that managers double down on after smaller losses. I demonstrate in panel regressions that managers avoid doubling the portfolio weights of losing positions. I also find tenuous evidence that managers facing more fund-level career risk, proxied by poor trailing manager-level returns, are particularly hesitant to substantially add to a losing position, relative to managers facing less career risk. While I cannot definitively prove that career risk is driving managers' hesitancy to double down, my results are consistent with this mechanism.

I construct a variety of control portfolios to demonstrate that "doubling down" is not explained by mechanical return effects or by previously identified asset pricing phenomena. In particular, my finding is not the result of a simple reversal effect, of a fund's best ideas (large positions), or of the general informativeness of fund trades. Funds exit, rather than double down on, most of the positions in which they suffer losses. The positions that they exit do not outperform. In my sample, funds' largest positions do not substantially outperform. Positions that managers double after strong trailing stock-level performance do not outperform to nearly the same extent as the double down positions. I also find that these "double up" positions fail to outperform the positions that a manager chooses to exit after strong trailing stock-level performance. In other words, large hedge fund manager trades are statistically informative only after poor trailing position-level performance. This finding is consistent with my proposed mechanism. For positions with strong trailing performance, managers are choosing between riding or harvesting winners, which has no clear implications for career risk.

Many studies that form portfolios that generate large outperformance figures construct hypothetical portfolios that face short sale constraints and high transaction costs. The double down portfolio instead represents actual positions of significant size held by hedge fund managers. These are long positions in stocks that are liquid enough for managers to make large trades in. The double down portfolio I construct at each quarter end is based, on average, upon nearly $2 billion of actual manager positions.

Doubling down represents a clear demonstration of the information content of a manager's portfolio management decisions, when she knows that her investors will be watching closely. These findings thus contribute to the literatures on career risk and selective manager skill. On the other hand, doubling down as I define it is rare in the context of the universe of all hedge fund equity positions. My results therefore have little to say about aggregate measures of skill.

My findings also suggest that during asset price dislocations, some of the specialists in those assets — fund managers that already own a stake in a given stock, for example — are hesitant to devote additional capital to those positions as a result of the losses they have already suffered. This reasoning extends the basic limits-to-arbitrage intuition (Shleifer and Vishny (1997)), which links fund-level performance to a manager's reluctance to take on additional risk. My findings are primarily the result of position-level, rather than fund-level, underperformance. I study position portfolio weights, rather than dollar position sizes, which should substantially reduce the impact of flows and past manager performance on my findings. These results add a new facet to the interaction between career risk and asset price dislocations.

2.1. Literature

While studies of aggregate skill in mutual funds and hedge funds have found mixed results, a persuasive literature has emerged that managers generate positive abnormal risk-adjusted returns on certain positions, which are identifiable ex-ante. For instance, hedge funds outperform on their confidential holdings (Agarwal, Jiang, Tang, and Yang (2013)), and mutual fund buys during large outflows (or sells during large inflows) are informative (Alexander, Cici, and Gibson (2006)). Relevant to my approach, several papers utilize mutual fund portoflio weights to predict outperformance, such as the best ideas of mutual funds (Cohen, Polk, and Silli (2010)), mutual fund active share (Cremers and Petajisto (2009)), and mutual fund industry concentration (Kacperczyk, Sialm, and Zheng (2005)). As a whole, this literature finds that manager conviction is at times related to performance, if one can properly identify a manager's strongest beliefs or expertise.

Lakonishok, Shleifer, Thaler, and Vishny (1991) detail the practice of window dressing by fund managers. They find that a sample of pension fund managers do tend to sell their losing positions, likely as an attempt to avoid investors making negative inferences of the managers' skill. When a manager doubles down, in contrast, the fund manager is adding to a losing position, thus calling even *more* attention to that position.

Working against my findings are two separate effects the literature has identified. First, mutual fund flows are known to chase past performance, and have furthermore been shown to predict future returns (Coval and Stafford (2007)). This effect should work against my findings, as managers double down on positions that have run against them. One would expect the poor past performance of double down positions to be associated with worse manager-level performance, *ceteris paribus*, and thus outflows. These outflows would in turn tend to drive negative future returns on those positions, as managers sell them to meet redemptions.

Second, the disposition effect in mutual funds appears to predict underreaction in asset prices (Frazzini (2006)). Doubling down is a bit different, as when a manager doubles down, they are not merely holding on to a loser, as the disposition effect would predict. Instead, managers are actually adding to losing positions. In the case of doubling down in hedge funds, I tend to find reversals in stock-level performance, rather than the drift that the disposition effect has been shown to predict. Managers only double down infrequently, however, so drift could still dominate in the full sample of losing positions. Furthermore, my focus on hedge funds, rather than mutual funds, may explain some of the differences in my findings.

On the theory side of things, Shleifer and Vishny (1997) formalize the seminal concept of limits to arbitrage. This argument provides an explanation for why rational arbitrageurs may limit the positions they take to correct mispricings. In their model, uninformed investors cause fund managers to face outflows following poor fund-level performance, making managers averse to taking on large amounts of risk. Scharfstein and Stein (1990) analyze why fund managers may be hesitant to deviate from popular positions as a result of relative performance evaluation metrics. While hedge fund manager contracts are not typically explicitly tied to the performance of other managers, anecdotally hedge fund investors evaluate funds relative to the available universe of hedge funds. When a manager doubles down on a position, she is likely to stand out from other hedge funds unless many other funds are doubling down in the same position (which I do not find empirically).

Savor and Gamboa-Cavazos (2011) illustrate that in the aggregate, short sellers cover their positions after losses (price increases, in their case). However, the effect they identify is at the aggregate short interest level, due to data limitations, and is thus fundamentally different from my own. Savor and Camboa-Cavazos are cleverly and effectively illustrating limits to arbitrage in the aggregate — that new arbitrageur capital is not fully replacing the losses and outflows of existing short sellers — rather than showing that existing short sellers are necessarily choosing to retreat. It is quite possible that fund managers are adding to their short positions as a percentage of their assets under management while aggregate short interest is declining as a result of the portfolio implications of shorting.[1] Furthermore, given

[1]For example, imagine stock X is worth \$100 per share. Suppose a single manager with \$100 in assets is short 1 share (100% of assets) of stock X. Suppose stock X increases to \$150. The manager will now have \$50 in assets, assuming no flows. Suppose the manager covers a third of a share of X. The manager will now be short \$100, or 2/3 of one share of stock X. The manager now has a 200% short position in X. The manager has increased her portfolio weight on her short position, but short interest in stock X has declined from 1 share to 2/3 of a share.

that outflows tend to follow poor past trailing performance, flow effects would also push towards a reduction in aggregate short interest after losses on short positions. I identify doubling down by examining managers' portfolio weights, and thereby limit the influence of changes in managers' total assets under management. I am thus able to more effectively infer individual manager beliefs regarding future returns.

I contribute to the selective skill literature by demonstrating that the path by which a fund manager reaches her portfolio weights can provide additional predictive power and larger magnitudes of inferred skill relative to most previously identified effects. On the other hand, doubling down is by its nature quite rare, and does not have broader implications for the total amount of skill in the institutional investing universe. My findings are consistent with a position-level career risks mechanism. I add to the traditional limits to arbitrage literature by focusing on this position-level, rather than fund-level, underperformance as another potential limitation faced by skilled investment managers trading against mispricings.

2.2. 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).[2] I impose simple backward looking filters to attempt to eliminate funds that file 13Fs that are clearly not representative of a manager's overall portfolio.

In more detail, 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.[3] The Form

[2] I thank the authors for kindly providing me with their hedge fund sample.

[3] More detailed requirements are provided at https://www.sec.gov/answers/form13f.htm. The full list of 13F

13F reports the list of 13F securities that the investment manager holds as of the end of the corresponding calendar quarter.[4] In panel regressions in Section 2.3, I employ the full sample, from 12/31/1980 through 12/31/2013. In later sections, I construct portfolios that represent positions held by funds from 12/31/1989 through 12/31/2013. I omit the beginning of the sample because the resulting portfolios are too thin during that period, when there are fewer observations.

I then filter for the 13Fs of hedge funds using the comprehensive list 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 hedge fund dataset.

I obtain stock returns 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 obtain data on DGTW returns from Russ Wermers' website.[5] Risk factor returns (SMB, HML, UMD) are from Ken French's website.

I assume that the set of securities filed on a fund's 13F constitutes a representative portfolio. I am trying to identify the potential expertise of active, "stockpicking" managers. Yet 13F filings do not provide information on short positions, cash holdings, or non-U.S. equity positions. I therefore remove filings that are clearly unrepresentative of a firm's investment strategy, or filings that identify firms pursuing strategies that are not likely to be based on active stockpicking. For example, a fund that reports only a single stock on a Form 13F is probably investing primarily outside of publicly listed U.S. equity holdings, while a fund that holds a controlling interest in a stock's common equity is likely pursuing a private

---

securities is available at http://www.sec.gov/divisions/investment/13flists.htm.

[4]As is common in the literature, in constructing my tests of performance, I ignore the 45 day filing delay. Instead, I analyze portfolios as of the date the manager holds the associated underlying positions. This approach focuses on the behavior of the managers themselves instead of attempting to construct a trading strategy that a third-party market participant could implement using only publicly available information.

[5]The DGTW benchmarks are available via http://www.smith.umd.edu/faculty/rwermers/ftpsite/Dgtw/coverpage.htm

equity strategy. Other funds, such as quant funds, hold a disproportionately large number of positions, and are less likely to rely on thorough stock-level analysis to make investment decisions. These funds are thus unlikely to have differential information on whether a change in the price of a stock is the result of a fundamental change in a firm's business prospects or whether the price change is simply noise, since they base their decisions primarily on aggregate patterns in accounting data and returns.

I therefore remove (1) any filing in which a single holding represents over 60% of the 13F portfolio, (2) any filing with fewer than 10 positions, (3) any filing in which a fund holds over 50% of the total outstanding market cap of a stock whose market cap exceeds $250 million, (4) any filing in which the value of the 13F portfolio is under $50 million, and (5) any filing which contains more than 150 positions. None of my results are sensitive to these particular threshold values. These filters reduce my sample of fund-quarters from 48,260 to 28,578.[6]

After imposing the filters above, the hedge fund 13F portfolios that remain should generally be representative of managers' beliefs within those portfolios. All discussion of manager returns and flows refer to these portfolios. If a manager holds a portfolio of 20 different stocks, in addition to cash, several short positions, and a number of credit positions, there is no reason that the manager's long stock position weights should not represent, on average, a manager's relative evaluation of different opportunities. For instance, if a manager overweights a given stock within her long portfolio, it seems reasonable to infer that the manager most likely believes that stock has a greater expected return or lower risk as compared to some of her other stock holdings.

Table 2.1 summarizes the hedge fund universe across 13F filings from 12/31/1989 through 9/30/2013. Averages are taken in the time series, with the datapoint in any given quarter

[6]The mutual fund literature employs comparable filtering techniques. See, for instance, Kacperczyk, Sialm, and Zheng (2007). The main differences between my approach and the standard mutual fund sample selection procedure are driven by the fact that mutual fund holdings data includes cash and non-stock holdings, and mutual funds are labeled with explicit investment objectives. Furthermore, mutual funds rarely take a private equity approach to investing.

representing an average that is equal-weighted across managers, but value weighted within any given manager's portfolio.[7]

Table 2.1. Hedge fund universe summary statistics

This table displays the characteristics of my hedge fund sample, after applying the filters described in the text. The sample covers 12/31/1989 - 9/30/2013. Statistics are taken across the full set of 96 13F filings covered in the sample, except for characteristic data which is across 90 13F filings (12/31/1989 - 3/31/2012). Quintile averages are weighted by portfolio weights. A value of 5 represents a higher measure of the underlying statistic, ie the largest market cap quintile, the highest book to market quintile, or the highest trailing 12-month performance (excluding the most recent month) quintile.

	Mean	Median	10th pctl	90th pctl	Std Dev
Hedge funds per quarter	293	259	77	530	174.7
Positions per quarter	17,011	17,131	5,409	28,004	8,566
Total long U.S. equity assets per quarter ($ BB)	219.9	151.9	31.4	493.1	180.1
Median position value ($ MM)	2.5	2.3	1.8	3.3	0.6
Avg position size quintile	3.9	3.9	3.7	4.1	0.2
Avg position book quintile	2.7	2.7	2.6	2.9	0.1
Avg position momentum quintile	3.2	3.2	2.9	3.4	0.2

Quintiles are from characteristic-based assignments (i.e., DGTW portfolio assignment). The sample of hedge funds tends to hold stocks in larger size quintiles, with above average momentum, and with slightly below average book-to-market. The sample grows steadily over time, and peaks at almost 600 managers in late 2007.

2.3. When do managers double the weight of a position?

I first examine what factors explain how managers alter their portfolio weights over time. If a position-level career risk mechanism is truly at work, then managers should be hesitant

[7]For example, these averages weight a 10% position of a manager with $1 billion in assets under management twice as heavily as a 5% position of a manager with $5 billion in assets under management.

to add to positions that have generated large past losses, relative to positions that have generated large gains. Furthermore, if position-level and fund-level career risks interact, it is possible that managers facing greater fund-level career risk may be hesitant to double down on losing positions, relative to managers facing less fund-level career risk.

Attempting to explain the complete panel of manager-stock-quarter portfolio weight changes would be a substantial task, given the vast universe of potential investments and the numerous sources of noise involved. Instead, I focus my attention on large increases in the portfolio weights of previously sizable positions. I do not attempt to explain small portfolio weight changes, or the portfolio weight changes of small positions. I also do not attempt to explain a manager's decision to initiate a new position in a particular security.

In other words, what factors lead a hedge fund manager to increase her bet on a position that she already holds? When do managers fulfill the "doubling" requirement of the phenomenon of "doubling down"? The dependent variable that I employ is an indicator variable that captures when a manager doubles a portfolio weight over the past 3, 6, 9, or 12 months. Doubling down, the details of which I explain in Section 2.4, occurs when a manager doubles the portfolio weight of a position specifically after poor trailing stock-level performance, a subset of the events captured by my indicator variable.

I first remove all positions that were not sizable as of time t-q from the sample (I set q=1, 2, 3, or 4). I define a sizable position as one with a portfolio weight greater than the maximum of (1) 2.5% and (2) a manager's average position size in all 13Fs she has filed to date (defined at the manager-quarter level).[8] I use this definition of a sizable position throughout the paper, although I later vary the exact cutoffs for robustness and to demonstrate comparative statics.

For the left hand side, I construct an indicator variable, $double_{s,m,t,q}$, that is set to 1

[8]I define average position size as the reciprocal of the average number of equity positions on all of a manager's 13Fs filed up to and including the date of analysis. For example, if a manager has filed 13Fs with 20, 30, and 40 positions to date, then her average position size will be $1/(\frac{1}{3} * (20 + 30 + 40)) = 3.33\%$.

if the position weight at time t of stock s for manager m at least doubled over the past q quarters, and 0 otherwise. As an example, suppose q=2. If Microsoft was a 5% position for manager m on June 30, 2004, and a 10% or greater position for manager m on December 31, 2004, then $double_{Microsoft,m,12/31/2004,2} = 1$. If, on the other hand, Microsoft was only a 6% position for manager m at December 31, 2004, then $double_{Microsoft,m,12/31/2004,2} = 0$.

I employ a linear regression approach with standard errors clustered in two dimensions, at the manager and quarter levels. While the dependent variable is an indicator variable, it is merely a simplification of the continuous trading behavior of managers that allows me to cleanly separate events. I am not overly concerned with small negative predicted values for the left hand side variable, which could loosely be interpreted as an increase in the likelihood of decreasing a position, rather than doubling its weight. My regressions never produce fitted values above 1. Since a manager's portfolio weights by definition add up to 1, a decision to double one position clearly impacts the decision to double other positions. Clustering errors at the manager level is therefore necessary. I additionally cluster by time, for robustness and in case portfolio weight decisions are correlated across managers during a given time period.

I employ three different regression frameworks. In all regressions, I include controls at the stock level for instititutional ownership as a percentage of total market capitalization (*IOR*), the number of institutional owners (*numInst*), the stock's book-to-market ratio (*BM*), and the log market capitalization (*logMktCap*). At the manager level, I include a control for a manager's 13F assets (*13Fassets*).

In my first regression, I include a variable that represents a stock's recent gain relative to the market (*recentGain*), expressed as a percentage contribution to the manager's portfolio. In other words, if stock s was a 5% position but outperformed the market by 20% over the past q quarters, then its *recentGain* would be positive 1% for that period. I winsorize returns relative to the market at the 1% level in this calculation. Losses are recorded as negative numbers.

$$double_{s,m,t,q} = \alpha_0 + \alpha_1 recentGain_{s,m,t,q} + \gamma' controls_{s,m,t,q} \tag{2.1}$$

My proposed position-level career risks mechanism would suggest that managers are averse to adding substantially to losing positions ($\alpha_1 > 0$). I strongly confirm this hypothesis in the data. In Panel A of Table 2.2, the coefficient on *recentGain* is positive and highly statistically significant (t-statistic of over 12 across all values of *q)*. A positive coefficient means that as *recentGain* becomes more negative (as losses increase in magnitude), a manager is less likely to double a position.

Table 2.2. Predicting manager doubles

This table displays the results of panel regressions to predict hedge fund manager position weight doubles. The dependent variable $double_{s,m,t,q}$ is an indicator variable set to 1 when a manager m doubles the portfolio weight of a position s at time t over the trailing q quarters (1, 2, 3, or 4, as denoted). *stockMinMkt* is the performance of a stock over the past q quarters relative to the CRSP value weighted market index. *previousPosSize* is the percentage weight of that position for manager m as of q quarters ago. *recentGain* is *stockMinMkt * previousPosSize*, or the gains/losses on that position over the past q quarters, relative to holding an equal-sized position in the market index. *trailingRetQuintile* is the quintile of the manager's 13F portfolio returns over the past q quarters, relative to the full sample of managers. *gainLess0* is an indicator set to 1 if *recentGain* is less than 0. Interactions are as denoted. Additional controls include a stock's institutional ownership percentage, its number of institutional holders, book-to-market, and log market capitalization, and the 13F assets of the fund manager. T-statistics are displayed in brackets, based on standard errors clustered by both manager and by time period. ** and * denote significance at the 5% and 10% levels, respectively.

Panel A: Impact of *recentGain*

Double position size over past Dependent variable:	3 months $double_{s,m,t,1}$	6 months $double_{s,m,t,2}$	9 months $double_{s,m,t,3}$	12 months $double_{s,m,t,4}$
recentGain	0.3665	0.4814	0.6069	0.6322
	[12.62]**	[14.09]**	[14.19]**	[15.05]**
IOR	0.0031	0.0046	0.0023	0.0042
	[1.71]*	[1.72]*	[0.64]	[1.05]
numInst	0.0000	0.0000	0.0000	0.0000
	[-5.03]**	[-3.46]**	[-1.61]	[-0.47]
BM	0.0025	0.0025	0.0012	0.0007
	[2.43]**	[2.27]**	[1.24]	[0.57]
logMktCap	0.0013	0.0002	-0.0014	-0.0027
	[3.58]**	[0.42]	[-2.05]**	[-3.39]**
13Fassets	0.0000	0.0000	0.0000	0.0000
	[-3.39]**	[-2.52]**	[-1.07]	[-1.43]
constant	-0.0183	0.0094	0.0500	0.0785
	[-2.56]**	[0.88]	[3.37]**	[4.61]**
Observations	247,407	202,539	168,488	141,473
R-squared	0.0021	0.0042	0.0075	0.0097

Of course, one interpretation of this coefficient is simply that managers are averse to trading positions, and instead choose primarily to let position weights drift based on per-

Table 2.2: (continued)

Panel B: Impact of *stockMinMkt* and *previousPosSize*

Double position size over past	3 months	6 months	9 months	12 months
Dependent variable:	$double_{s,m,t,1}$	$double_{s,m,t,2}$	$double_{s,m,t,3}$	$double_{s,m,t,4}$
stockMinMkt	0.0323	0.0419	0.0511	0.0540
	[14.19]**	[15.44]**	[16.13]**	[16.89]**
previousPosSize	-0.0753	-0.1166	-0.1330	-0.1563
	[-10.16]**	[-10.27]**	[-9.81]**	[-9.33]**
IOR	0.0032	0.0042	0.0010	0.0021
	[1.68]*	[1.49]	[0.28]	[0.49]
numInst	0.0000	0.0000	0.0000	0.0000
	[-4.26]**	[-2.27]**	[-0.04]	[1.38]
BM	0.0026	0.0030	0.0023	0.0026
	[2.63]**	[2.72]**	[2.37]**	[2.33]**
logMktCap	0.0010	-0.0005	-0.0026	-0.0041
	[2.66]**	[-0.94]	[-3.56]**	[-4.95]**
13Fassets	0.0000	0.0000	0.0000	0.0000
	[-3.05]**	[-2.22]**	[-0.76]	[-1.05]
constant	-0.0083	0.0301	0.0802	0.1165
	[-1.15]	[2.72]	[5.21]**	[6.44]**
Observations	247,407	202,539	168,488	141,473
R-squared	0.0041	0.0082	0.0139	0.0183

Table 2.2: (continued)

Panel C: Interactions

Double position size over past	3 months	6 months	9 months	12 months
Dependent variable:	$double_{s,m,t,1}$	$double_{s,m,t,2}$	$double_{s,m,t,3}$	$double_{s,m,t,4}$
*recentGain*gainLess0*	-0.0364	-0.0469	-0.0208	-0.0344
**trailingRetQuintile*	[-1.20]	[-1.64]	[-0.80]	[-1.38]
recentGain	0.0211	-0.0473	0.0341	-0.1355
	[0.29]	[-0.71]	[0.43]	[-1.62]
stockMinMkt	0.0533	0.0681	0.0794	0.0850
	[13.55]**	[14.27]**	[15.34]**	[15.58]**
previousPosSize	-0.1135	-0.1752	-0.2140	-0.2284
	[-12.15]**	[-13.79]**	[-13.93]**	[-12.94]**
*trailingRetQuintile*recentGain*	0.0103	0.0191	-0.0219	-0.0049
	[0.42]	[0.73]	[-1.05]	[-0.24]
*trailingRetQuintile*gainLess0*	0.0004	-0.0005	-0.0011	-0.0007
	[1.49]	[-1.06]	[-1.95]*	[-1.53]
*recentGain*gainLess0*	-0.5351	-0.5769	-0.7220	-0.5271
	[-5.20]**	[-7.42]**	[-7.88]**	[-5.16]**
gainLess0	0.0033	0.0082	0.0121	0.0128
	[4.33]**	[5.18]**	[6.62]**	[8.02]**
trailingRetQuintile	-0.0009	-0.0003	-0.0002	-0.0003
	[-3.13]**	[-0.67]	[-0.42]	[-0.56]
IOR	0.0039	0.0050	0.0024	0.0033
	[2.19]**	[1.81]*	[0.67]	[0.80]
numInst	0.0000	0.0000	0.0000	0.0000
	[-4.48]**	[-2.39]**	[-0.29]	[1.12]
BM	0.0027	0.0031	0.0023	0.0025
	[2.78]**	[2.90]**	[2.46]**	[2.31]**
logMktCap	0.0012	-0.0001	-0.0019	-0.0035
	[3.35]**	[-0.15]	[-2.64]**	[-4.10]**
13Fassets	0.0000	0.0000	0.0000	0.0000
	[-2.56]**	[-1.96]*	[-0.42]	[-0.88]
constant	-0.0151	0.0170	0.0599	0.0959
	[-2.08]**	[1.56]	[3.91]**	[5.20]**
Observations	247,407	202,539	168,488	141,473
R-squared	0.0050	0.0097	0.0163	0.0208

formance. However, hedge funds are known to have relatively high turnover. Agarwal, Fos, and Jiang (2013) find that hedge funds turn over their 13F portfolios about 0.92 times a year on average. Furthermore, I am not analyzing high frequency trades. Over horizons of 9 or 12 months, it seems unlikely that a manager is especially averse to altering her portfolio weights, since she will typically turn over almost her entire portfolio over such a horizon. The strong significance of the coefficient on *recentGain* over these longer horizons — with higher magnitudes and t-statistics than over 3 and 6 month horizons — provides some reassurance that portfolio weight drift is not the dominant cause of managers doubling positions.

In Panel B, I break down this effect into a stock's previous position size as a percent of a manager's 13F portfolio as of quarter *t-q* (*previousPosSize*), and the performance of the stock relative to the market over the past q quarters, winsorized at the 1% level (*stkMinMkt*).

$$double_{s,m,t,q} = \alpha_0 + \alpha_1 stkMinMkt_{s,m,t,q} + \alpha_2 previousPosSize_{s,m,t,q} + \gamma' controls_{s,m,t,q} \quad (2.2)$$

Managers appear less likely to add to stocks that have poor market-adjusted performance, based on the positive and significant coefficient on *stkMinMkt* (α_1). They are also less likely to double large positions ($\alpha_2 < 0$); this makes sense from a portfolio management perspective, if one values diversification and does not want to allow a single position to dominate the portfolio.

Finally, in Panel C I undertake a more ambitious approach that uses a proxy for fund-level career risk, the quintile of a manager's trailing 2 year return (*trailingRetQuintile*).[9] A manager with returns in the top quintile over the past 2 years is likely more secure from a career perspective than a manager in the bottom return quintile. My hypothesis is that fund-level career risk has a particular impact on a manager's willingness to double a position after losses, relative to after gains. I therefore also include an indicator (*gainLess0*) set to

[9] Returns are measured as the buy and hold returns of a manager's 13F portfolio, rebalanced at each quarter end to the manager's latest 13F filing. Quintiles at each date are based on the trailing returns of all fund managers in my sample at that date with sufficient return history. A higher quintile represents higher trailing returns.

1 when *recentGain* is less than 0. I include the three-way interaction and all two-way interactions between *recentGain, gainLess0,* and *trailingRetQuintile.* The coefficient of interest, β, is the coefficient on the three-way interaction between these variables. I retain all other variables used so far as RHS variables.

$$\begin{aligned} double_{s,m,t,q} &= \alpha_0 + \beta recentGain_{s,m,t,q} * gainLess0_{s,m,t,q} * trailingRetQuintile_{s,m,t,q} \\ &+ \delta' twoWayInteractions_{s,m,t,q} + \alpha_1 recentGain + \alpha_2 stkMinMkt_{s,m,t,q} \\ &+ \alpha_3 previousPosSize_{s,m,t,q} + \gamma' controls_{s,m,t,q} \end{aligned} \tag{2.3}$$

In other words, β is the coefficient on the interaction term between specifically a stock's past losses (rather than gains) and a proxy for career risk (a manager's trailing 2 year return quintile). I have separately included the other interactions, so tests of β examine whether or not career risks have a differential impact on a manager's willingness to double down after *losses* in a position, *relative* to after gains.

The consistently negative estimate of β across all specifications suggests that a manager facing less career risk (*trailingRetQuintile* is more positive) who is holding a position that has generated more losses (*recentGain* is more negative) may be more willing to double the weight of that position than is a manager facing more career risk.[10] This effect is relative to the unconditional relationship between career risk and gains/losses. However, statistical significance is weak. One regression just misses the 10% significance threshold (t-statistic 1.64 for q=2) for β, while the others are not as close. There is thus only marginal evidence that the impact of fund-level career risk may be different for gains than for losses in the direction I have hypothesized.

Overall, these results suggest that managers are reluctant to double the portfolio weight of a position on which they have lost considerable money. There is also tenuous evidence that managers facing elevated career risk at the fund level are particularly reluctant to double a

[10] More losses*less career risk*coefficient = negative*position*negative = positive.

losing position.

2.4. Double down portfolio construction

In this section, I define doubling down in the data. I then provide summary statistics of the double down portfolio.

2.4.1. Definition

In order to test my hypothesis, I translate the anecdotal story that I described in the introduction into a quantifiable procedure. I define doubling down as a stock s held by a manager m at the end of quarter t that meets the following criteria over the past q quarters (I refer to q as the portfolio formation period). First, the position s must have been sizable for manager m in period t-q.[11] Second, over the last q quarters, stock s's return must have fallen short of the CRSP value weighed index return by more than $Z\%$. Third, the manager must have increased the weight of s in her portfolio at time t to at least G * (the position weight of s at time t-q).

I hold a stock in the double down portfolio until the stock is no longer a sizable position for manager m. In other words, I hold the position until the "investment thesis plays out," from the manager's perspective, or until the manager exits the sample (whichever comes first). I weight each position using its portfolio weight for manager m, divided by the sum of the portfolio weights of all positions in the portfolio (so that position weights sum to 1). I utilize this approach because it allows relative position sizes within manager portfolios to matter, but ignores the size of managers' total U.S. equity portfolios.[12] My results carry

[11] As explained in Section 2.3, I define a sizable position as one with a portfolio weight greater than the maximum of (1) 2.5% and (2) a manager's average position size in all 13Fs she has filed to date (defined at the manager-quarter level). Later, I will vary these cutoffs for robustness, i.e., I use 3% instead of 2.5%, or I use 1.5 times a manager's average position size.

[12] There is some empirical (see, for example, Cremers and Petajisto (2009)) evidence that skill declines among

through on an equal weighted-basis.

Finally, I exclude instances in which a manager doubles down on a stock over the past q quarters when the proportional change in the manager's assets over those q quarters was in the bottom decile of all the managers in my sample. In other words, I ignore observations when a manager doubles down over a period during which her 13F assets dropped precipitously. I remove these observations because they are much less likely to reflect a manager's beliefs regarding future excess returns. Managers that have experienced a rapid drop in their 13F assets are typically facing large outflows, have suddenly shifted their assets towards other (non-U.S. equity) strategies, or have suffered extreme negative returns and will soon face large outflows in the future. In the first two instances, if a manager has rapidly reallocated her funds to cash or to other asset classes, it is likely that liquidity considerations played a large role in determining portfolio weight changes. Similarly, in the third case, the flows-performance literature documents that if a manager has suffered extreme negative returns, the manager will likely face large outflows in the future. Such a manager will need to raise cash in anticipation of future outflows, once again meaning that liquidity considerations will have an outsized impact on portfolio weights. Furthermore, my portfolio construction approach relies on a manager's decision to reduce a position's weight to determine when to remove that position from the double down portfolio. Knowing that a manager may soon be forced to liquidate means that there is a significant chance that I will be forced to remove any associated positions when the manager liquidates, rather than waiting for an information based sell signal for that position. In other words, I will likely be unable to follow such a stock until the manager thinks it is no longer undervalued. Finally, I also remove these observations of doubling down to differentiate my effect from Alexander, Cici, and Gibson (2006), who find that manager purchases during extreme outflows generate

---

the largest managers.

future outperformance.[13]

In my baseline portfolio construction, I employ the following parameter values. I test q=2 quarters, Z=10%, and G=2. I later proceed to vary all of these parameters (q=1-4, Z=0%-15%, G=1.5-2.25) for robustness, and to demonstrate that outperformance increases when doubling down is conditioned on greater past losses.

2.4.2. Summary statistics

Table 2.3 summarizes the double down portfolio formed using the baseline parameter values listed above. Since I examine managers' sizable positions, the average size quintile of the double down portfolio increases relative to the full sample, as one would expect. The average book-to-market quintile of these stocks is below that of the median stock (the third quintile). The average momentum quintile is similarly below that of the median stock, as one would expect given the portfolio formation procedure.

Figure 2.1 displays the number of double down positions in my portfolio over time (left hand scale) against the number of hedge funds in my sample (right hand scale). The number of double down positions grows in line with the number of hedge funds. The average total value of these underlying positions in their respective managers' portfolios, across the 96 quarters in the sample, is \$1.84 billion.[14] Even the lowest 10th percentile of the value of the portfolio is over \$180 million. In other words, the double down portfolio represents actual substantial bets made in the market by the underlying hedge fund managers in my sample. Taking another perspective, managers separately double down on 410 positions, and hold these positions for between 4 and 5 quarters (13 months), on average.

[13] Appendix Table B.1 provides baseline performance results for the double down portfolio without imposing this filter. The point estimates are slightly lower, but are still quite statistically significant.

[14] At each quarter end, I add up the dollar values, in their respective managers' 13F filings, of all of the double down positions. For instance, if the portfolio has two positions, and hedge fund X holds \$600 million of one position while hedge fund Y holds \$300 million of the second, I would record a value of \$900 million.

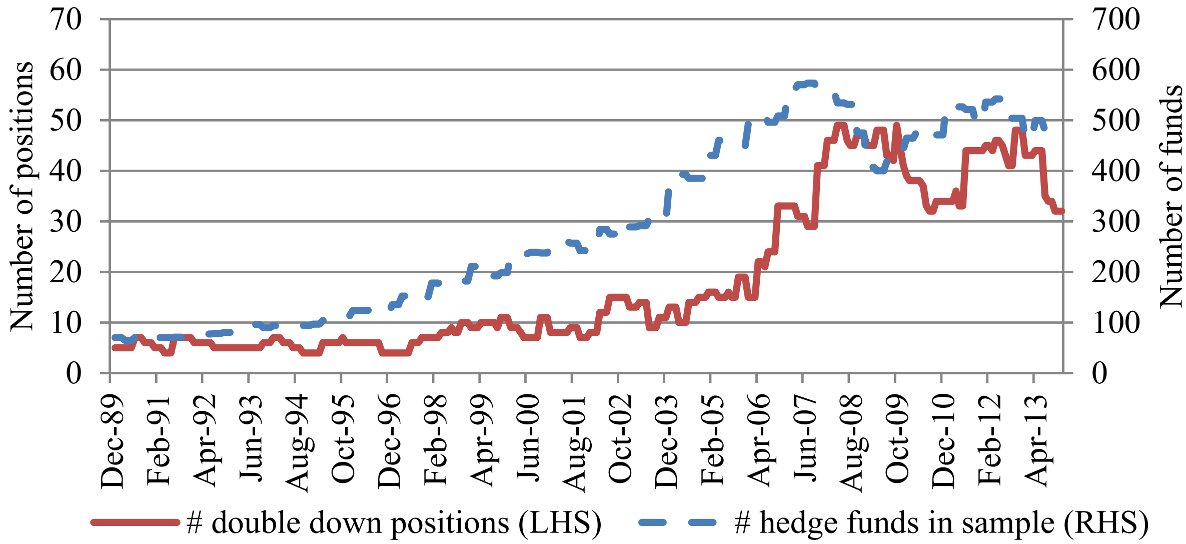

Figure 2.1. Double Down Portfolio Composition over Time
This figure displays the number of positions in the double down portfolio (left hand scale, solid line), constructed as described in the text, against the number of hedge funds in the sample (right hand scale, dashed line).

2.4.3. Exiting instead of doubling down

The double down portfolio constructed above suggests that doubling down is a rare occurence, as one would expect if managers take on substantial career risk by doing so. To provide some context for the frequency of doubling down, beyond the panel regressions in Section 2.3, I compare the number of double down positions to the number of positions that managers *exit*, rather than double down on. I use this portfolio again in Section 2.6 as a control portfolio for performance tests.

I construct an exit portfolio of positions that are double down eligible, but which managers halve the position weight of, as opposed to double, over the relevant time frame. In other words, these positions meet all other requirements to enter the double down portfolio, except the requisite increase in position weight, which instead declines substantially. I then hold each such position for 4 quarters after the manager has exited it, since this is

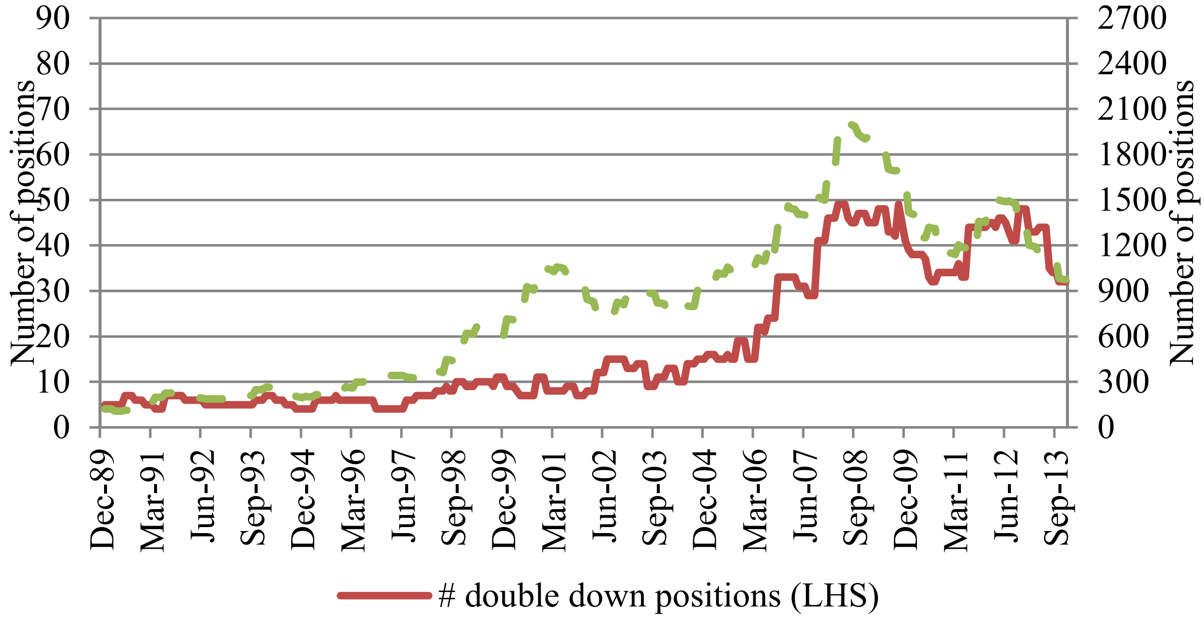

Figure 2.2. Double Down or Exit?
This figure displays the number of positions in the double down portfolio (left hand scale, solid line), constructed as described in the text, against the number of positions that managers exit instead of double down on (right hand scale, dashed line).

approximately the average holding time of positions in the double down portfolio. Figure 2.2 displays the number of stocks in the double down portfolio (left hand scale) compared to the number of stocks in a portfolio of positions that managers have chosen to exit, rather than double down on (right hand scale). As the difference in scales makes evident, managers exit roughly 30 positions for each position that they double down on. As another comparison, on average across the sample period, the set of all hedge funds in aggregate holds roughly 200 sizable non-double down positions for each double down position.

2.5. Double down portfolio performance

I have demonstrated that managers are more willing to double the portfolio weight of a position that has done well recently than they are willing to double the portfolio weight of a position that has done poorly. This finding suggests that managers are hesitant to add to losing positions. However, arguably a better test of whether managers are more selective when adding to losing positions is to examine the performance of the losing positions to which managers add. If managers have a higher threshold for adding to losing positions, as a result of the ensuing career risk, then one would expect to see that when managers actually do make such a decision, those positions outperform.

In this section, I test the risk-adjusted performance of the double down portfolio. I demonstrate that managers do indeed outperform on those holdings. I then extend these performance tests to focus on comparative statics, and find results consistent with a career risks mechanism. That is, I demonstrate that doubling down after greater position-level losses predicts greater future outperformance.

2.5.1. Performance - baseline

Table 2.4 displays risk-adjusted performance measures of the double down portfolio. The 4-factor alpha is significant at the 5% level when doubling down occurs over a 3, 6, or 9 month interval. DGTW-adjusted performance and CAPM-alphas are significant for doubling down over 6 and 9 months. These figures are strongly positive and close to significant for the 3 month portfolio. At 6 or 9 months, outperformance figures range from 48 to 83 bps per month. Annualized, those figures correspond to outperformance of 5.8% (12*.0039) to 10.0% (12*.0083). At 3 months, monthly outperformance figures range from 39 bps to 78 bps.

At 12 months, all point estimates remain positive, but the doubling down effect begins to break down, as none of the outperformance estimates are statistically significant. It should not be too surprising that the effect of doubling down dissipates when using a sufficiently long portfolio formation horizon. The premise of the doubling down mechanism I propose is that a manager increases her portfolio weight in a stock in response to its underperformance. Hedge funds are highly active investors. It is not surprising that a manager takes less than

Table 2.3. Double down portfolio summary statistics

This table displays the characteristics of the double down portfolio, formed as described in the text using a 6-month portfolio formation lookback window. The sample covers 12/31/1989 - 9/30/2013. Statistics are taken across the full set of 96 13F filings covered in the sample, except for characteristic data which is across 90 13F filings (12/31/1989 - 3/31/2012). Quintile averages are weighted by portfolio weights. A value of 5 represents a higher measure of the underlying statistic, ie the largest market cap quintile, the highest book to market quintile, or the highest trailing 12-month performance (excluding the most recent month) quintile.

	Mean	Median	10th pctl	90th pctl	Std Dev
Hedge funds per quarter	17	10	4	38	13.8
Positions per quarter	18	10	5	44	15.3
Total long U.S. equity assets per quarter ($ BB)	1.84	0.82	0.18	5.36	2.0
Median position value ($ MM)	39.7	35.9	22.1	54.1	25.5
Average position size quintile**	4.2	4.2	3.5	5.0	0.5
Average position book quintile**	2.5	2.6	1.5	3.2	0.6
Average position momentum quintile**	2.8	2.7	2.1	3.6	0.7

12 months to respond to a potential buying opportunity after a stock drops in price.

On the other hand, the effect is statistically a bit weaker at 3 months than at 6 or 9 months. This finding is driven by the fact that there are fewer observations of doubling down at the 3 month horizon. There are fewer stocks that managers double in position size and that have fallen a full 10% short of the market in such a short time frame. Given the size of these positions, for instance, holding all else equal (such as assets under management), it may take many weeks for a manager to fully double her portfolio weight in a stock. If the manager is trading in reaction to a drop in the price of a stock, to have both the price drop and the subsequent portfolio weight change occur within a single calendar quarter, which contains roughly 60 trading days, is apparently less common than to have both these events occur over the course of 2 or 3 quarters.

Naturally, with more frequent data on hedge fund holdings and trades, I would be able to more precisely observe both the time horizons at which doubling down operates and how managers time their trades. The outperformance of the double down portfolio using

Table 2.4. Double down portfolio

This table displays the monthly performance of the double down portfolio, formed as described in the text. The baseline parameter values are used here. That is, over the relevant portfolio formation period, a stock's return must fall short of the CRSP value weighted market index by at least 10%, and the manager must have increased the position portfolio weight to 2 * its weight at the begining of the formation period. Furthermore, the position must be sizable at both the beginning and end of the formation period, with sizable defined at the manager quarter level as the maximum of (1) 2.5% and (2) the manager's average position size across all 13Fs filed by the manager to date. Portfolio performance is calculated from 12/31/1989-12/31/2013 for alpha calculations, and from 12/31/1989-6/30/2012 for DGTW calculations. Positions are weighted equally across managers but value-weighted within a given manager's portfolio, as described in the text. T-statstics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively.

Trailing ret interval	Raw return	DGTW adjusted	4 factor alpha	mkt	size	book	mom	CAPM alpha	mkt	Avg # positions in port†
3 mo	1.33%	0.39%	0.78%	1.09	-0.04	0.27	-0.31	0.55%	1.24	9.6
		[1.15]	[2.13]**	[11.5]	-[0.3]	[1.8]	-[2.7]	[1.54]	[15.8]	
6 mo	1.69%	0.83%	0.72%	1.14	0.07	0.35	-0.18	0.64%	1.24	18.3
		[3.70]**	[3.35]**	[17.9]	[0.7]	[3.4]	-[3.8]	[2.84]**	[20.7]	
9 mo	1.41%	0.63%	0.49%	1.13	0.18	0.35	-0.15	0.48%	1.21	20.0
		[2.66]**	[2.03]**	[17.8]	[1.7]	[3.2]	-[2.0]	[1.99]**	[19.4]	
12 mo	1.22%	0.35%	0.24%	1.11	0.26	0.36	-0.12	0.28%	1.16	26.6
		[1.63]	[1.12]	[19.0]	[2.6]	[5.4]	-[2.0]	[1.26]	[20.0]	

†The 3-month trailing return portfolio has no positions for 2 of the 96 quarters in my sample. The 6-, 9-, and 12- month portfolios are populated for all 96 quarters.

quarterly observations, which is still based on large changes in manager portfolio weights, suggests that managers are able to time their trades at least reasonably well when doubling down. If managers do in fact have some timing ability in this particular circumstance — when adding to large positions that they have lost money on — then one might expect that this timing ability would show up more cleanly in weekly or monthly holdings data, generating even larger outperformance figures.

The portfolio by construction weights negatively on momentum, or the UMD coefficient. The positive weight on value, or HML, of the double down portfolio is only slightly larger than the HML loading on the full hedge fund sample. The portfolio weights slightly above unity on the market, and generally a small positive amount on size.

Figure 2.3 displays the trailing 3-year CAPM and 4-factor alphas of the double down portfolio using a 6 month (Panel A) and 9 month (Panel B) formation period. As is apparent, outperformance is not generated solely during a small subsample, nor is it generated only early in the sample when the portfolio is based on fewer underlying positions. 3-year trailing alphas are very rarely substantially negative.

Going forward, I focus on the double down portfolios formed at the 3, 6, and 9 month horizons, with particular emphasis on results at the 6 and 9 month horizons, where I have more observations of doubling down.

2.5.2. Performance - comparative statics

If the career risk mechanism detailed so far is truly driving my results, one would expect that the outperformance of the double down portfolio would increase as its portfolio formation cutoffs are tightened. In other words, if a manager has suffered greater past losses on a position, then the act of doubling down on that position would be expected to generate more career risk for the manager. In turn, one would expect the double down positions formed conditioning on greater past position-level losses to have higher expected returns in

Panel A: 6 month portfolio formation period

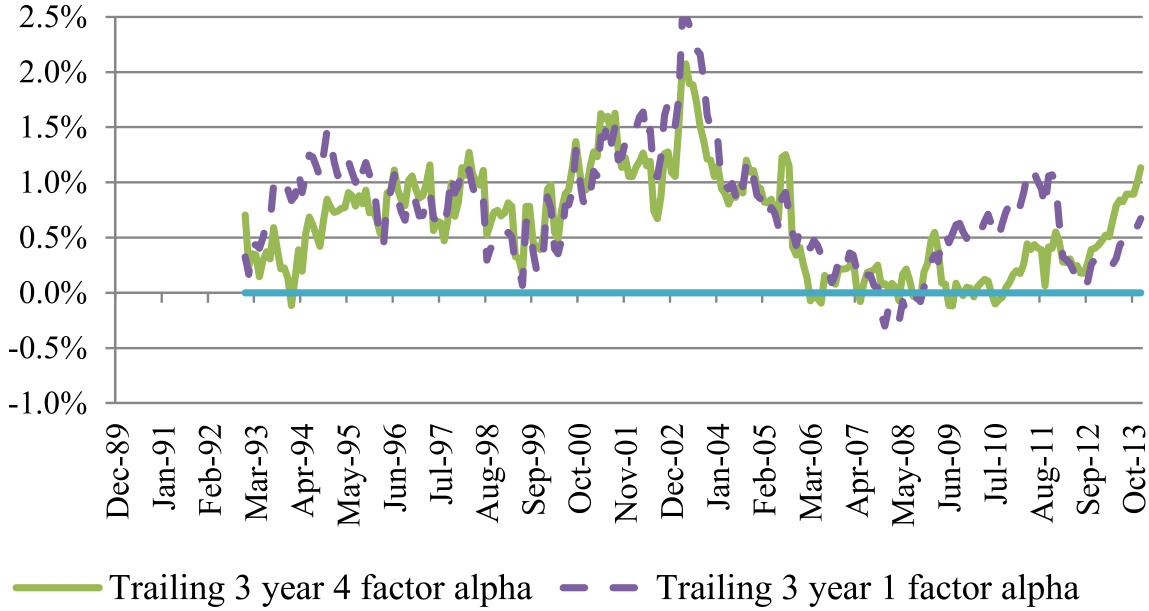

Panel B: 9 month portfolio formation period

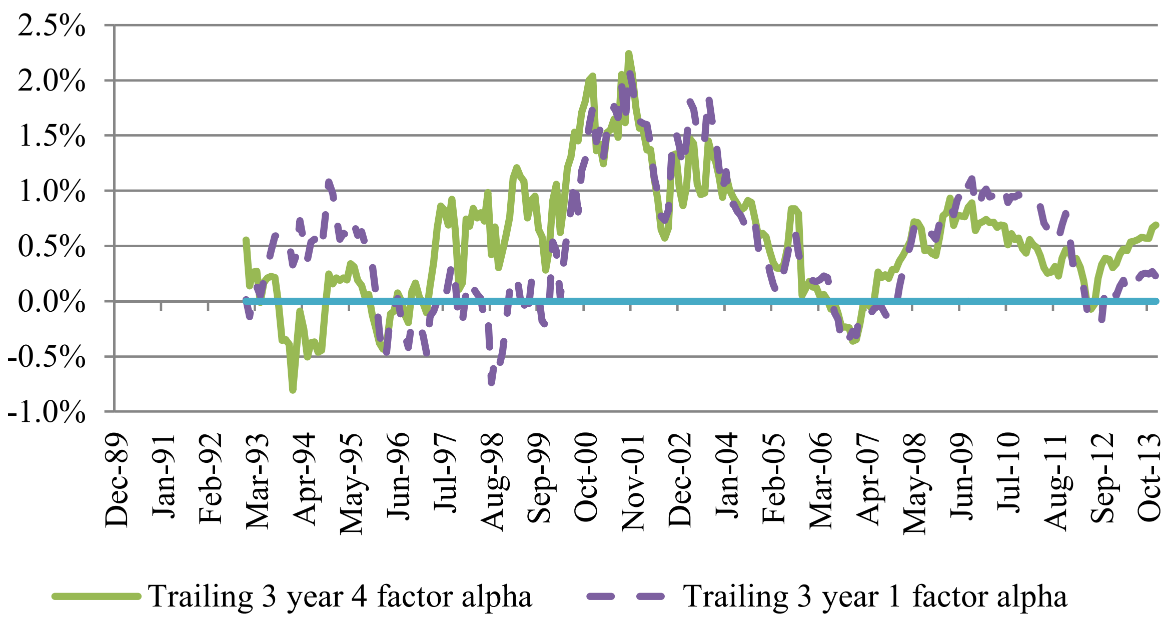

Figure 2.3. Double Down Portfolio 3-year Trailing Performance
This figure displays the monthly trailing 3-year 4-factor (solid line) and CAPM alpha (dashed line) of the double down portfolio, constructed as described in the text.

order to offset this career risk. Varying the portfolio formation parameters in this manner also provides a robustness check against data snooping.

I confirm this hypothesis in the data. I obtain risk-adjusted performance estimates as high as 123 bps per month, which leads to annualized outperformance of roughly 15%. For brevity, I only display 4-factor alphas for the double down portfolio formed using a 6 month window in Table 2.5. The Appendix displays the full tables (B.2-B.5) of DGTW-adjusted performance figures, 4-factor alphas, and CAPM alphas of portfolios formed using 3, 6, and 9 month windows with varying cutoffs. The full results are similar.

Tightening the double down cutoffs by definition reduces the number of double down observations. The resulting portfolios are therefore thinner than my baseline portfolios. On the other hand, though point estimates decrease, the statistical significance of my result remains even when I loosen the cutoffs relative to my baseline parameters. These portfolios of course are comprised of a greater number of underlying positions than the portfolios in my baseline construction.

Increasing the sizable position cutoff by varying either the absolute (2.5%) or the relative (the multiple of a manager's average position size) floor means that a manager will be required to have lost more in a given position, for the same stock-level performance, prior to doubling down in that position. As expected, as I increase the minimum sizable portfolio weight from 1.5% to 3.5%, the four-factor alpha increases from 41 bps per month to 106 bps per month. Varying the relative position cutoff, outperformance increases from 62 bps per month when a cutoff of 0.5 times a manager's average position size is used to 129 bps per month when a position must have started at 2 times a manager's average position size to be considered sizable.

Conditioning on different past performance cutoffs (the parameter Z) also varies a manager's past losses in a given position. Greater past losses generate greater outperformance after doubling down. A portfolio formed of the stocks that fell short of the market by 5%

Table 2.5. Comparative Statics

This table displays the monthly performance of the double down portfolio, formed as described in the text, but varying the parameter values used to form the portfolio. In each column, a single parameter value (X) is varied, as displayed, relative to the baseline case. The resulting 4-factor alpha of the 6-month formation period portfolio is displayed. In the first column, the definition of sizable is the maximum of (1) X% and (2) the manager's average position size across all 13Fs filed by the manager to date. In the second column, the definition of sizable is the maximum of (1) 2.5% and (2) X * the manager's average position size across all 13Fs filed by the manager to date. In the third column, over the past 6 months, a stock's return must fall short of the CRSP value weighted market index by at least X%. In the fourth column, over the past 6 months, the manager must have increased the position portfolio weight to X * its weight at the begining of the formation period. Portfolio performance is calculated from 12/31/1989-12/31/2013. Positions are weighted equally across managers but value-weighted within a given manager's portfolio, as described in the text. T-statstics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively.

Varying position cutoff		Varying average position factor		Varying fall relative to market		Varying increase in portfolio wtd	
New Value	Char.-adj monthly perf	New Value	Char.-adj monthly perf	New Value	Char.-adj monthly perf	New Value	Char.-adj monthly perf
1.50%	0.41%	0.50	0.62%	0%	0.38%	1.50	0.29%
	[2.57]**		[2.96]**		[2.06]**		[1.87]*
2.00%	0.57%	0.75	0.59%	5%	0.46%	1.75	0.47%
	[2.99]**		[2.87]**		[2.28]**		[2.51]**
3.00%	0.81%	1.50	0.94%	15%	1.01%	2.25	1.20%
	[2.95]**		[3.16]**		[3.30]**		[3.55]**
3.50%	1.06%	2.00	1.29%				
	[2.88]**		[3.64]**				

over the last 6 months, which a manager doubles down on, generates a 4-factor alpha of 46 bps per month. On the other hand, a portfolio of double down positions that fell short of the market by 15% generates a 4-factor alpha of 101 bps per month.

Finally, I consider different requisite increases in a manager's portfolio weight in a position (the parameter G) to define doubling down. One would expect that the larger you make a position after its past drop, the more salient it will be to investors. Once again, risk-adjusted performance increases from 29 bps to 120 bps per month as G goes from 1.5 to 2.25.

Double down positions generate greater outperformance, on average, when they are initiated following larger position-level losses.

2.5.3. Performance - event study

The calendar time portfolio is the preferred statistical test of the outperformance of a portfolio. However, examining all double down positions pooled together is a method of checking the robustness of the portfolio approach. In particular, one might be concerned that because the portfolio approach equally weights some quarters with few double down positions and other quarters with a large number of double down positions, it could be producing misleading results. The fact that the trailing 3-year alpha of the portfolio is not systematically different across most subsamples suggests that this issue should not be a major concern. I take an event study style approach here to further demonstrate robustness.

I construct an event study approach by treating a manager's decision to double down in a position as a single event. I then equal weight across all events. In my first approach, I throw away all information about how managers trade these positions after they double down. I define doubling down using my baseline parameters. Date 0 is the portfolio formation date, the date at which each double down event occured. I DGTW-adjust the performance of each stock in the portfolio.

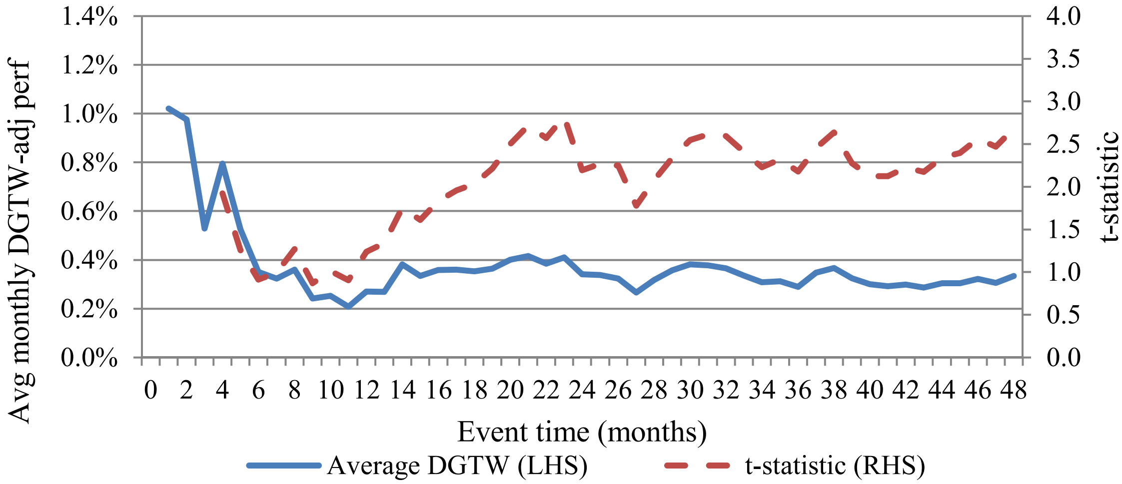

Figure 2.4. Event Study, Hold Regardless of Subsequent Manager Activity
This figure displays the equal-weighted average DGTW-adjusted performance to date (left hand scale, solid line) and corresponding t-statistic (right hand scale, dashed line) of the pool of all double down positions in event time. This figure treats date 0 as the date in which a manager doubles down in a position. In this figure, calculations are made by holding positions regardless of a manager's trading behavior subsequent to doubling down.

Figure 2.4 displays the results. Performance is displayed over time as the average of monthly DGTW-adjusted portfolio performance figures from date 0 until the corresponding date in event time. The t-statistic is generated based on this series of monthly DGTW-adjusted returns. Remarkably, I find consistent outperformance at long horizons here. A manager doubling down in a stock reliably predicts positive outperformance of that stock over the next four years, on the order of 30-40 bps a month, on a DGTW-adjusted basis. The weakest time frame of outperformance is over the 3-5 quarter horizon.

Of course, if a mispricing corrects soon after a manager doubles down on a position, then observing a manager's decision to exit that position is important. I thus also consider an approach that equal weights all positions in the double down portfolio, but exits a position when a manager does. For example, the portfolio generating performance from 18 to 21 months in this approach is the equal-weighted average DGTW-adjusted performance of all double down positions that managers continued to hold (as sizable positions) as of the

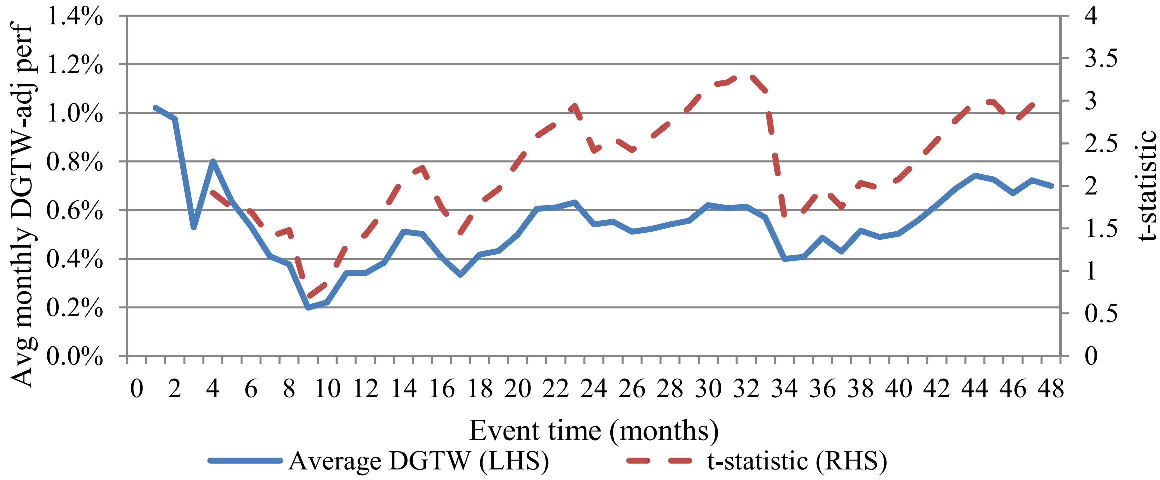

Figure 2.5. Event Study, Remove When Manager Sells
This figure displays the equal-weighted average DGTW-adjusted performance to date (left hand scale, solid line) and corresponding t-statistic (right hand scale, dashed line) of the pool of all double down positions in event time. This figure treats date 0 as the date in which a manager doubles down in a position. In this figure, calculations are made by removing positions from the underlying portfolio when a manager sells that position.

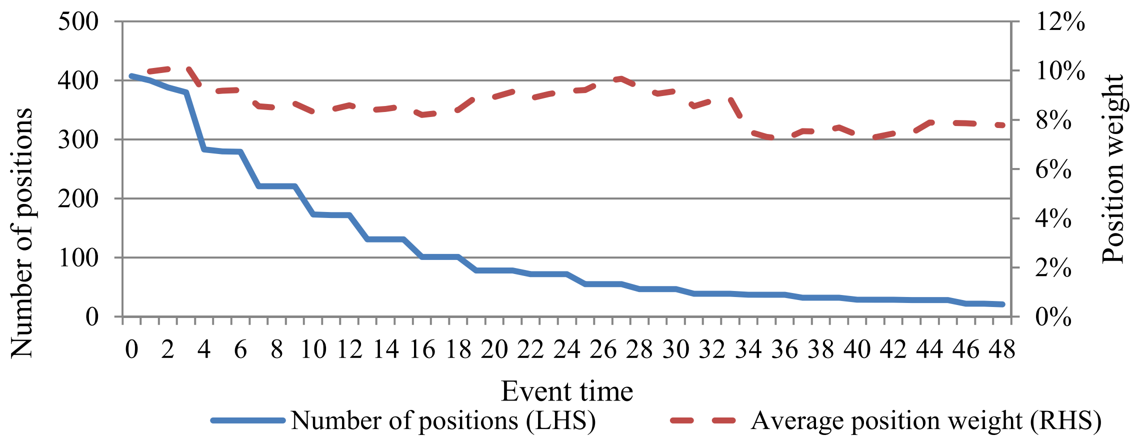

Figure 2.6. How Managers Exit and Size Double Down Positions
This figure displays how long managers hold each of their double down positions (left hand scale, solid line), treating date 0 as the date in which the manager doubled down in that position. It also displays the average position size of the remaining double down positions (right hand scale, dashed line).

6th quarter end following the quarter end at which they originally doubled down on those positions.

Figure 2.5 displays these results. Outperformance is now much stronger, though it does still dip briefly in the medium run. At longer horizons, the portfolio generates DGTW-adjusted outperformance estimates of between 40 and 80 bps, roughly in line with my portfolio results. Furthermore, it is worth noting that the performance of both of these portfolios is very strong in the first 2 quarters following the double down event. Since managers exit many positions after 1 or 2 quarters, this performance is weighted more heavily in the performance generated by the calendar time double down portfolio in Section 2.5.1.[15]

To illustrate this point, Figure 2.6 displays how managers exit and size their double down positions over time. The average holding time of a double down position is 13 months. As can be seen, managers exit about one third of these positions within the first two quarters after doubling down. By the fifth quarter, they have exited another third. Managers slowly exit the remaining positions over time. Managers size the positions that are in the double down portfolio to be between 8% and 10% of their 13F assets, on average.

2.6. Control portfolios

In this section, I illustrate that the outperformance of the double down portfolio is robust to controlling for a number of alternative explanations. Double down positions outperform other large hedge fund positions and positions that are double down eligible but which managers

[15] To illustrate this point, suppose that a fund holds two positions at all times across 40 quarters. One is a position in the same stock X for the entire sample. The other position is a different stock every quarter. Suppose the portfolio always weights both positions equally. Each quarter, the portfolio's performance will be generated as one half of the performance of stock X, and one half of the performance of the other stock in the portfolio. However, in constructing a pooled event approach such as I do in this section, using the purchase of a position as the date-0 event, the performance of the event portfolio would be based upon 40 positions in the first quarter, but thereafter only a single position for the remaining time in the event-study. This construction would downplay the performance contribution of the rotating positions in the portfolio, relative to their contribution to the performance of the actual fund over time.

instead exit. Other positions that managers buy on a dip do not outperform. Positions with large portfolio weight increases following strong position-level returns do generate some outperformance, though much less than their double down counterparts. Furthermore, these positions are statistically indistinguishable from positions with strong trailing returns that managers choose to exit rather than double. There seems to be less information content in a manager's portfolio weight changes after strong position-level performance, relative to following poor position-level performance. One would expect this differential if a manager devotes more attention to her portfolio management decisions regarding positions with poor performance, because of their potential career implications.

2.6.1. Full hedge fund sample

Table 2.6 displays the performance of the full sample of hedge fund equity positions, using the same weighting scheme as the double down portfolio. It also displays the performance of managers' largest positions (what one might expect to be their "best ideas," Cohen, Polk, and Silli (2010)).

Interestingly, I find that the managers in my filtered sample do generate economically small but statistically significant outperformance. Risk-adjusted performance ranges from roughly 10 to 20 bps monthly for the full set of hedge fund positions. These figures could potentially cover a management fee of between 1% and 2%, though they would have more trouble covering both a management fee and an incentive fee. Of course, managers could also potentially add (or subtract) value on the short side of their portfolios, which I am unable to observe, and could add (or subtract) value in non-U.S. equity positions or from intra-quarter trading.

The evidence for best ideas is much weaker in my hedge fund sample than in the mutual fund sample of Cohen, Polk, and Silli (2010). Sizable positions in my sample, defined in Section 2.3, generate performance in line with the full sample of fund positions. The perfor-

Table 2.6. Hedge fund sample, all and large positions

This table displays the monthly performance of the full sample of hedge fund positions and the subset of large hedge fund positions. Sizable is defined at the manager quarter level as the maximum of (1) 2.5% and (2) the manager's average position size across all 13Fs filed by the manager to date. Top position and top 3 positions are the portfolios formed of the single largest or largest three positions, by portfolio weight, in each manager's 13F portfolio at each quarter end. Portfolio performance is calculated from 12/31/1989-12/31/2013 for alpha calculations, and from 12/31/1989-6/30/2012 for DGTW calculations. Positions are weighted equally across managers but value-weighted within a given manager's portfolio, as described in the text. T-statstics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively.

	Raw return	DGTW adjusted	4 factor alpha	market	size	book	mom	CAPM alpha	market
all	1.01%	0.11%	0.15%	1.05	0.04	0.27	-0.02	0.18%	1.10
positions		[3.50]**	[3.13]**	[68.3]	[1.6]	[16.6]	-[1.6]	[2.66]**	[63.0]
sizable	1.01%	0.10%	0.15%	1.05	0.02	0.23	0.03	0.20%	1.08
positions		[2.77]**	[2.90]**	[64.4]	[0.9]	[12.9]	[2.2]	[3.08]**	[58.9]
top	1.08%	0.09%	0.25%	1.07	-0.08	0.25	0.10	0.16%	1.10
position		[1.22]	[3.16]**	[50.3]	-[2.4]	[7.5]	[4.6]	[2.41]**	[62.3]
top 3	1.05%	0.10%	0.17%	1.06	-0.01	0.24	0.08	0.16%	1.11
positions		[1.87]*	[2.61]**	[55.2]	-[0.3]	[10.3]	[4.8]	[2.29]**	[62.0]

mance of funds' single largest or top 3 positions is similar. The outperformance of double down positions, a subset of large positions, cannot be attributed to this effect.

To further illustrate this point, I go long the double down portfolio and short the complete set of sizable hedge fund positions. I display the results in Table 2.7, using double down portfolios formed after a stock dips relative to the market by 10% or 15%. The estimated long-short portfolio alphas are all strongly positive, and many are statistically significant.

2.6.2. Other positions with poor recent stock-level performance

I have demonstrated that if a manager doubles down on a stock after it has underperformed, the stock tends to do well going forward. But what if this is merely some sort of mechanical reversal effect? What happens to the stocks that underperform, but which managers do not

Table 2.7. Long double, short other large positions

This table displays the monthly performance of long-short portfolios that go long the double down portfolio and short the set of all sizable hedge fund positions. The double down portfolio is constructed as described in the text, using baseline parameter values but requiring that over the relevant portfolio formation period a stock's return must fall short of the CRSP value weighted market index by either 10% or 15%, as noted. Sizable is defined at the manager-quarter level as the maximum of (1) 2.5% and (2) the manager's average position size across all 13Fs filed by the manager to date. Portfolio performance is calculated from 12/31/1989-12/31/2013 for alpha calculations, and from 12/31/1989-6/30/2012 for DGTW calculations. Positions are weighted equally across managers but value-weighted within a given manager's portfolio, as described in the text. T-statstics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively.

	4-factor alpha			CAPM alpha		
	3 mo	6 mo	9 mo	3 mo	6 mo	9 mo
Long double down (<mkt-10%),	0.63%	0.57%	0.35%	0.35%	0.44%	0.28%
short other large positions	[1.74]*	[2.74]**	[1.49]	[0.98]	[2.02]**	[1.21]
Long double down (<mkt-15%),	1.25%	0.86%	0.33%	1.01%	0.79%	0.30%
short other large positions	[2.48]**	[2.87]**	[1.22]	[2.04]**	[2.55]**	[1.10]

double down on?

Table 2.8 displays the performance of relevant control portfolios that are formed conditional on poor recent trailing stock-level returns. The dominant effect that concerns the continuation of short term returns, of course, is momentum. A position's performance over the past year tends to be positively correlated with its performance over the following year. As I will show, none of the control portfolios I form here generate excess performance unless they are given credit for a (predictably) large negative weight on UMD. Relative to a DGTW-adjusted or market benchmark, these portfolios generate insignificant performance. Claiming to beat only a strongly short-momentum portfolio is rarely an objective of hedge fund managers. The double down portfolio, on the other hand, does well on both a DGTW-adjused and CAPM basis. Regardless, none of these comparable portfolios come close to matching the magnitude of the outperformance of the double down portfolio.

The first portfolio is constructed of positions that are double down eligible, but which managers exit (i.e., cut the position weight in half over the relevant time frame) rather

Table 2.8. Other trades after poor performance

This table displays the monthly performance of control portfolios formed of stocks that have recently underperformed the market. That is, over the relevant portfolio formation period, these stocks have performance that falls short of the CRSP value weighted market index by 10% or more. In the first control portfolio, an exit is a position that is double down eligible but which a manager cuts the portfolio weight of by half rather than doubling it. In the second portfolio, dip positions are those that managers initiate for the first time after poor trailing performance. In the third portfolio, large dip positions are positions that managers make sizable after poor trailing performance. Portfolio performance is calculated from 12/31/1989-12/31/2013 for alpha calculations, and from 12/31/1989-6/30/2012 for DGTW calculations. Positions are weighted as described in the text. T-statstics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively.

	Dgtw-adjusted			4-factor alpha			CAPM alpha		
	3 mo	6 mo	9 mo	3 mo	6 mo	9 mo	3 mo	6 mo	9 mo
Exit instead of double down	-0.05%	0.03%	0.03%	0.04%	0.17%	0.19%	-0.27%	-0.17%	-0.13%
	-[0.33]	[0.19]	[0.23]	[0.43]	[1.68]*	[1.98]**	-[1.62]	-[0.95]	-[0.72]
Buy any position on a dip	0.05%	0.08%	0.10%	0.10%	0.23%	0.20%	-0.09%	0.01%	-0.03%
	[0.47]	[0.90]	[1.05]	[1.02]	[2.54]**	[2.32]**	-[0.72]	[0.10]	-[0.21]
Buy large position on a dip	0.19%	0.15%	0.05%	0.25%	0.37%	0.22%	-0.07%	-0.06%	-0.20%
	[1.13]	[0.93]	[0.34]	[1.61]	[2.66]**	[1.97]**	-[0.38]	-[0.30]	-[0.99]

than double down on, as in Section 2.4.3. In other words, these positions meet all other requirements to enter the double down portfolio, except the requisite increase in position weight (which declines, rather than increases).[16] I then hold each such position for 12 months after the manager has exited it. I equal weight positions both within and across managers.[17] Results are similar for portfolios that are value weighted by a stock's market capitalization. DGTW and CAPM alphas are near zero or somewhat negative.

The second portfolio is formed of all positions that managers buy into on a dip (i.e., stocks with returns over the portfolio formation period that fall short of the market by 10%, and which managers did not hold at the start of the formation period). As with the double down portfolio, I hold all such positions as long as the fund manager does. Once again, these stocks do not generate enviable performance.

The third portfolio specifically looks at when managers initiate a sizable position after poor trailing performance. In other words, a position may or may not have been a small position before, but the manager increased its weight across the sizable threshold following stock-level underperformance. I hold all such positions as long as the manager continues to hold the stock as a sizable position. These positions display similarly meager performance.

In summary, it does not appear that fund positions generally outperform following stock-

[16] I specifically examine effective exits, rather than all positions that are double down eligible but which managers do not double down on, to make the control portfolio comparable to the double down portfolio. Requisite to doubling down, a manager must make a large change in the portfolio weight of a position. One likely reason for such a change is if a manager changes her beliefs regarding a position's degree of underpricing. Inversely, to exit a position, a manager must also make a large change in the portfolio weight of a position, potentially expressing the manager's belief that the position is no longer as attractive an investment. On the other hand, when a manager makes only a small change to the portfolio weight of a position that has underperformed recently, it is difficult to infer if the manager has changed her beliefs, or if other sources of noise such as transaction costs are driving the small portfolio weight change. Nevertheless, in unreported tests, available upon request, I construct portfolios of the positions that are double down eligible but which managers neither exit nor double down on. They generate performance results in line with the full set of sizable positions.

[17] I need to make some assumptions here because the manager has already effectively exited the position. I therefore cannot infer how long to hold the position, or how to weight it, from the manager's trading behavior, as I do in the construction of the double down portfolio. Managers hold double down positions for roughly 4 quarters on average, so I hold these stocks for a similar time period.

level dips.

2.6.3. Large portfolio weight changes following strong stock-level performance

What if a manager doubles her position weight in a stock after strong, rather than weak, trailing stock-level performance? Conditional on a manager doubling her portfolio weight in a stock, how does the stock's recent trailing performance matter?

If one thinks that large changes in a manager's portfolio weights are informative in general, one would expect that positions which managers increase in weight following stock-level outperformance would also do well. Notably, however, the fund flows literature would already lead us to expect this to be the case. If a fund holds a sizable position in a stock that does well, then the fund will tend to be doing better than average, and will attract inflows because flows chase performance. Flows will drive the price of the stock upwards, as that fund and other managers who held that stock receive inflows and buy more of their existing positions. This effect is of course correlated with momentum, though they are not the same. Coval and Stafford (2007), for instance, find that a portfolio formed of stocks held by funds with expected future inflows generates a positive monthly 4-factor alpha of about 30 bps. Controlling for momentum therefore only partially removes the effect of fund flows.

Table 2.9 displays the performance of positions that managers adjust after strong trailing performance (a return greater than the market plus 10% over the relevant time frame).[18] "Double up" positions that managers double after strong performance do well, though not nearly as well as their double down counterparts. On a 4-factor alpha or CAPM-alpha basis they generate outperformance of between 12 and 40 bps (in the vicinity of the Coval and Stafford estimates for inflow-based outperformance), compared to a range of 48-72 bps for

[18] I construct this portfolio just as I do the double down portfolio, in that positions are weighted by their manager-level portfolio weight, and are held as long as the manager continues to hold the stock as a sizable position.

Table 2.9. Trades after good performance

This table displays the monthly performance of control portfolios formed of stocks that have recently outperformed the market. That is, over the relevant portfolio formation period, these stocks have performance that exceeds the CRSP value weighted market index by 10% or more. In the first control portfolio, the manager doubles the portfolio weight of the selected positions. In the second portfolio, the manager exits, or cuts the portfolio weight of in half, the selected positions. Portfolio performance is calculated from 12/31/1989-12/31/2013 for alpha calculations, and from 12/31/1989-6/30/2012 for DGTW calculations. Positions are weighted as described in the text. T-statstics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively.

	Dgtw-adjusted			4-factor alpha			CAPM alpha		
	3 mo	6 mo	9 mo	3 mo	6 mo	9 mo	3 mo	6 mo	9 mo
Double after	0.25%	0.48%	0.30%	0.12%	0.41%	0.30%	0.22%	0.40%	0.33%
strong perf	[0.83]	[2.93]**	[1.72]*	[0.36]	[2.40]**	[1.70]*	[0.66]	[2.11]**	[1.71]*
Exit after	0.24%	0.21%	0.11%	0.14%	0.22%	0.13%	0.20%	0.30%	0.21%
strong perf	[2.40]**	[2.88]**	[1.56]	[1.34]	[2.59]**	[1.45]	[1.58]	[2.78]**	[2.03]**

the double down portfolios formed over the same formation horizons.

I demonstrated in Section 2.6.2 that double down eligible positions which managers exit instead of doubling down on do not outperform. In other words, managers have some skill at selecting which positions to double down on; changes in a manager's portfolio weight in a position after poor stock-level performance are informative. Analogously, what about positions that managers exit (cut the position weight of in half) instead of double up on after strong stock-level outperformance? The second portfolio in Table 2.9 shows that these positions generate 4-factor alphas or CAPM-alphas of between 13 and 30 bps, just under the outperformance generated by the double up portfolio.[19] In other words, the large positions that a manager holds and that have done well in the past tend to do well in the future, even after adjusting for momentum, regardless of whether the manager increases or decreases her portfolio weights in those positions.

To formally test this claim, in Table 2.10 I form several long-short portfolios. First, I

[19] I construct this portfolio just as I do the portfolio of positions that managers exit instead of douling down on. Positions are equal-weighted and are held for 12 months.

Table 2.10. Long double, short exit by trailing performance

This table displays the monthly performance of long-short portfolios that examine the informativeness of manager trades after differing levels of trailing position-level performance. In particular, these portfolios go long positions that managers double the portfolio weight of, and short positions that managers cut the portfolio weight of by half. The trailing stock-return requirement for each portfolio, relative to the market, is denoted in the table. The first portfolio looks at positions that, over the relevant portfolio formation period, have performance that falls short of the CRSP value weighted market index by 10% or more. The second portfolio looks at positions that beat the index by 10% or more. The third portfolio looks at positions that fall in between. Portfolio performance is calculated from 12/31/1989-12/31/2013. Positions are weighted as described in the text. T-statstics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively.

	4-factor alpha			CAPM alpha		
	3 mo	6 mo	9 mo	3 mo	6 mo	9 mo
After weak perf	0.79%	0.60%	0.32%	0.94%	0.91%	0.66%
(double down; <mkt-10%)	[2.19]**	[2.79]**	[1.29]	[2.73]**	[3.89]**	[2.66]**
After strong perf	-0.02%	0.19%	0.17%	0.02%	0.11%	0.12%
(>mkt+10%)	-[0.07]	[1.08]	[0.97]	[0.07]	[0.63]	[0.72]
After medium perf	0.08%	0.28%	0.14%	0.09%	0.35%	0.21%
(mkt-10%<r<mkt+10%)	[0.51]	[1.86]*	[0.54]	[0.53]	[2.37]**	[0.79]

go long the double down portfolio, and short the positions that a manager exits instead of doubling down on. I then go long the double up portfolio, and short the positions that a manager exits instead of doubling up on. Finally, I go long positions that a manager doubles following middling performance (stock-level performance within 10% above or below the market over the relevant time frame), and short those that a manager exits after similar recent performance.

The double down long-short portfolio generates strong outperformance. In contrast, the double up long-short portfolio generates point estimates of alphas below 20 bps that are statistically indistinguishable from zero at all portfolio formation horizons. The final portfolio formed of stocks with middling recent returns generates performance figures between those of the other two portfolios.

Fund manager portfolio weight changes after strong trailing position-level performance are not informative. As expected, positions that have done well in the past continue to outperform, regardless of a manager's actions.

2.6.4. Adjusting for industry performance

What if managers are merely rotating into down-and-out industries at the right time, such as buying financial stocks in early 2009? To test this concern, I industry adjust the performance of my double down portfolio using the performance of each position's respective industry under the Fama-French 48 industry classification system. Table 2.11 displays the results. The double down portfolio generates significant outperformance in excess of an industry matched portfolio. Point estimates diminish marginally relative to a DGTW adjustment. Going long the double down portfolio and short each stock's respective industry portfolio generates positive and significant 4-factor alphas using 3, 6, or 9 month portfolio formation periods. Doubling down does not simply represent timely industry bets.

Table 2.11. Double down portfolio, industry adjusted

This table displays the monthly industry-adjusted performance of the double down portfolio, and the monthly performance of long-short portfolios that go long the double down portfolio and short an industry-matched portfolio. The double down portfolio uses baseline parameters. Industry returns are generated by matching each position in the double down portfolio to its corresponding Fama-French 48 industry. Portfolio performance is calculated from 12/31/1989-12/31/2013. Positions are weighted equally across managers but value-weighted within a given manager's portfolio, as described in the text. T-statstics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively.

Industry-adjusted			4-factor alpha of long double down, short industry			CAPM alpha of long double down, short industry		
3 mo	6 mo	9 mo	3 mo	6 mo	9 mo	3 mo	6 mo	9 mo
0.20%	0.75%	0.49%	0.78%	0.72%	0.49%	0.55%	0.64%	0.48%
[0.59]	[3.39]**	[2.06]**	[2.13]**	[3.35]**	[2.03]**	[1.54]	[2.84]**	[1.99]**

2.6.5. Active portfolio weight changes

In a frictionless setting, managers set their portfolio weights at every date based upon their beliefs regarding future asset returns. Of course, frictions such as transaction costs and taxes mean that managers do not trade as frequently as the frictionless benchmark would imply, a consideration I discussed briefly in Section 2.3. Yet hedge funds have relatively high turnover and I am not analyzing high frequency trades. Furthermore, I am only looking at large portfolio weight changes. Departing from the frictionless benchmark does not seem to be warranted in my case.

Nonetheless, to illustrate the robustness of the doubling down effect, I construct a control portfolio that utilizes specifically a fund's *active* portfolio weight changes. An active portfolio weight change over the past q quarters is the change in the position's portfolio weight minus the change that would have occured in the position's portfolio weight if the manager had made absolutely no trades over the quarter. In other words, $activeweightchange = currentpositionweight - (1+returnofstock)/(1+returnofmanagerm)*previouspositionweight$, with all returns measured over the portfolio formation period.

In this case, I define doubling down as before, except that I now require a manager to

Table 2.12. Double down portfolio, active weight changes

This table displays the monthly performance of the double down portfolio, formed as described in the text, but requiring that over the relevant portfolio formation period, the manager must have actively increased the position portfolio weight by at least the position's size at the beginning of the formation period. Otherwise, the baseline parameter values are used. Portfolio performance is calculated from 12/31/1989-12/31/2013 for alpha calculations, and from 12/31/1989-6/30/2012 for DGTW calculations. Positions are weighted equally across managers but value-weighted within a given manager's portfolio, as described in the text. T-statstics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively.

Trailing ret interval	Raw return	DGTW adjusted	4 factor alpha	mkt	size	book	mom	CAPM alpha	mkt
3 months	1.35%	0.44%	0.82%	1.18	0.01	0.16	-0.27	0.64%	1.29
		[1.54]	[2.80]**	[12.5]	[0.1]	[1.3]	-[2.8]	[2.15]**	[17.3]
6 months	1.48%	0.65%	0.54%	1.17	0.23	0.30	-0.21	0.50%	1.24
		[3.00]**	[2.65]**	[19.5]	[2.4]	[3.3]	-[4.8]	[2.26]**	[22.1]
9 months	1.26%	0.46%	0.43%	1.09	0.25	0.30	-0.23	0.38%	1.17
		[2.22]**	[1.94]*	[18.7]	[2.7]	[3.1]	-[3.7]	[1.64]	[19.0]

actively increase the portfolio weight of a position over the past q quarters by the position's original portfolio weight at t-q. For example, suppose stock s starts as a 5% position at time t-q. The stock declines by 50% over the next two quarters, but the overall fund has a return of 0%. If the manager did not trade any positions over those q quarters, then stock s would now be a 2.5% position. In order for the manager to actively increase her position in stock s by its original portfolio weight, stock s will have to be a 7.5% or larger position at time t.

Table 2.12 displays the performance of double down portfolios that are formed using active weight changes over 3, 6, and 9 month portfolio formation periods. The double down portfolio continues to generate strong outperformance. For instance, point estimates decline only slightly at 9 months. This result makes sense, as empirically, trading frictions are not sufficient to prevent managers from trading extensively over this time horizon (as is made clear by the fact that managers on average turn over almost their entire portfolio every 12 months).

2.7. Conclusion

Hedge fund managers outperform substantially and significantly on the positions that they double down on. Portfolios formed of these positions generate risk-adjusted outperformance of 5-15% on an annualized basis. The outperformance of doubling down is not explained by mechanical or previously identified asset pricing effects. Rather, doubling down behavior and returns are consistent with a career risks mechanism for this effect. Doubling down on a stock reverses the phenomenon of window dressing. By adding to a loser, fund managers call more attention to their mistakes.

If managers are hesitant to add to losing positions, then this effect may limit the amount of arbitrage capital that trades against mispricings. Existing holders likely represent a significant portion of the group of specialists who understand an individual asset well enough to separate mispricings from fundamental underperformance over short time horizons. If existing holders are constrained, and available specialist capital is limited, then past position losses will have implications for the dynamics of asset price dislocations beyond the well-known impact of fund-level losses. This possibility warrants further investigation.

3. Stock Experts

P.A.W. Capital Partners, a hedge fund based in Greenwich, CT, originally purchased Microsoft stock in the late-1990s. By December 31, 1997, P.A.W. held 5% of the fund's portfolio in Microsoft. In order to justify such a large position, P.A.W. likely researched the company extensively. P.A.W. proceeded to sell its stake in Microsoft prior to the most extreme throes of the technology bubble. During the quarter ending June 30, 2005, however, P.A.W. once again saw something that it liked in the stock, and added Microsoft back into its portfolio. Microsoft had underperformed recently, falling short of the S&P 500 index over the prior 6 months. P.A.W.'s investment proved profitable – over the next seven quarters, Microsoft returned 31%, beating its characteristic-matched portfolio by 33% and the S&P 500 by 16%.

The case of P.A.W. and Microsoft reflects a recurring phenomenon. In any body of knowledge, participants develop expertise about a subset of potential topics, and rely on that knowledge in their future work. In economics, after learning the intracacies of a dataset or the technical conditions behind an equilibrium model in order to write a paper, economists return to this acquired expertise to produce additional research. In the world of fundamental investment research, fund managers specialize in learning about particular companies within their universe of potential investments. After learning the details of a company's business model, its key competitors, and its critical suppliers, fund managers leverage this expertise in the future by reestablishing positions in stocks that they have held in the past.[1]

[1] An alternative interpretation is that fund managers create large positions based upon preexisting networks

In this paper, I study the expertise that investment funds develop in specific companies by focusing on the long U.S. equity portfolios of hedge fund managers, which are filed with the SEC under form 13F. To distinguish expertise from superficial knowledge, I focus on positions that are a significant percentage of a fund's portfolio. It is difficult to infer at exactly what point in time a fund becomes an expert in a given stock, however. A fund could establish a large position in a stock prior to being an expert in that company if, for example, the fund manager believes that the stock is sufficiently undervalued. Alternatively, a fund could thoroughly research a company and slowly build a position over time as it gradually develops and fine tunes a discounted cash-flow model of the company's business. Thus the first time that a fund purchases a given stock does not provide a well-identified event to analyze. On the other hand, one would certainly expect that by the time a manager has exited a large position, she has become an expert in that stock. Assuming that expertise decays sufficiently slowly, then if a manager subsequently chooses to *reenter* such a stock, that purchase reflects a manager's decision to trade in a stock in which she is an expert.[2]

I find that fund managers reenter a disproportionate number of the stocks that they hold over the fund's observable lifetime. To fix ideas, I define this concept of an "expert position" by a sequence of three events. In chronological order, first, a fund must have established a large position in a given stock in the past. Second, the fund must have effectively exited its position in that stock. Third, the fund must subsequently make that stock a large position in its portfolio. Upon completion of this sequence, that holding becomes an expert position until the fund manager exits the stock. I find that fund managers establish expert positions in a disproportionate number of the stocks that they have made into large positions in the

---

through which fund managers glean insider information, and that these networks, rather than the knowledge gained, is persistent (see, for example, Cohen, Frazzini, and Malloy (2008)). While the source of this form of specialization is different, its implications are similar. It is also possible that both forms of specialization coexist.

[2] I rely on buy and sell decisions as important signals because the literature documents that these portfolio shifts are stronger signals of conviction than simply examining fund holdings. See, for example, Rhinesmith (2016) or Alexander, Cici and Gibson (2007).

past. In particular, taking the set of all of a fund manager's large positions, I find that a manager will establish expert positions in 8.9% of these stocks, on average. Using simulation methods, I demonstrate that the odds of observing this proportion is nearly zero if a manager was equally likely to invest in any listed equity at a given point in time.[3] This finding suggests that managers believe they derive some value from their expertise (or else they would not keep coming back to the same stocks).

If managers know these stocks better than others, what would we expect to observe when a manager decides to establish an expert position? On the one hand, a manager might establish an expert position because she expects it to generate risk-adjusted outperformance. She may be quicker to perceive a mispricing than other managers as a result of her preexisting knowledge of the company. On the other hand, even when a manager does not expect a stock to outperform, she might establish an expert position as a portfolio placeholder. If the manager does not have better opportunities, then such an "active" investment would more effectively justify the fund's management fee than an equivalent investment in a passive index or cash. At the same time, the manager would be able to leverage her expertise when discussing the position with the fund's investors – the manager would have to do relatively more work on a non-expert stock to discuss it intelligibly.

Empirically, I find that trailing stock-level performance helps distinguish between these two scenarios. When a manager establishes an expert position in a stock that has underperformed the median stock over the past several quarters, it is a signal that the stock's price drop is not warranted and that the stock will outperform on a risk-adjusted basis going forward. The corresponding portfolio outperforms by 5-10% annualized. Two natural control portfolios fail to match this performance. First, managers do not outperform on the set of non-expert stocks that they purchase after poor recent stock-level returns. Second, a

[3]That null is the appropriate one if, for example, the underlying private-information generating process and the cost of acquiring that information is the same for all stocks.

characteristic-matched portfolio of the set of all stocks with poor recent stock-level returns fails to outperform. This latter portfolio demonstrates that general asset price reversals do not explain my results. On expert positions initiated after strong recent performance, however, managers do not outperform, consistent with the "portfolio placeholder" hypothesis.

Finally, I examine how the amount of available expert capital plays a role in price dynamics incremental to the effect of current ownership. I find that available expert capital significantly predicts stock-level volatility. That is, I find that the amount of available expert capital in a stock, normalized by the stock's market cap, significantly predicts the standard deviation of daily or biweekly returns during the subsequent quarter. This finding is robust to using a Fama-Macbeth regression with numerous controls, or in a panel setting with stock-level fixed effects and standard errors clustered by time. Experts in a stock are likely to be actively monitoring the underlying company's performance even if they do not currently hold that stock. Their presence may lead market makers to increase spreads to account for an increased expected proportion of informed trades. On the other hand, available expert capital may pick up the non-information based actions of other market participants. Distinguishing between these two conceptual explanations is a topic for future work.

Overall, this evidence highlights the relevance of stock-level expertise among hedge funds. It seems natural that a similar concept might apply to the investments of hedge funds in other asset classes. The network of investment expertise and relationships may affect capital formation: if a fund has developed expertise in a given firm, it may be a prime candidate for trading against a firesale or providing refinancing when credit is tight. Thus if the network matters, then the level of losses to the fund management industry will not be a sufficient statistic for predicting the real impact of a disruption: the distribution of losses across funds matters, too.

The paper proceeds as follows. In Section 3.1, I review related literature. In Section

3.2, I outline my data sources and data sample. In Section 3.3, I show that managers have a tendency to come back to their expert stocks and I construct test portfolios. In Section 3.4, I find that available expert capital predicts return volatility. I include select robustness results in Section 3.5. I conclude in Section 3.6.

3.1. Literature

The behavior of stock experts is naturally connected to manager skill. While studies of aggregate skill in mutual funds and hedge funds have found mixed results, a persuasive literature has emerged that managers generate positive abnormal risk-adjusted returns on certain positions, which are identifiable ex-ante. Agarwal, Jiang, Tang, and Yang (2013), for instance, show that hedge funds strongly outperform on their "confidential holdings." Most relevant for my work, several studies have demonstrated that managers outperform on the positions on which they might be expected to have the most information. Rhinesmith (2016) shows that funds outperform in stocks in which they purchase a large share of the total trading volume. Cohen and Polk (2010) find that a portfolio of mutual fund managers' "best ideas" – fund managers' largest risk-adjusted positions – generates excess risk-adjusted returns. Cohen, Frazzini, and Malloy (2008) demonstrate the outperformance of mutual funds on stocks in which the portfolio manager shares educational ties to the company's board. Kazperzyk, Sialm, and Zheng (2005) find that mutual fund managers with portfolios concentrated in fewer industries (and who presumably therefore devote more attention per industry) outperform other managers. These studies collectively support the idea that managers exhibit skill on at least some of the positions we might expect them to have more information on ex-ante. None of these studies, however, examines the persistence of expertise, or the tendency of managers to trade in and out of the positions that they have developed an expertise in.

On the theoretical side, the work of van Niewerburgh and Veldkamp (2009) and Merton (1987) provides a basis for why managers would want to concentrate their portfolios in their most specialized positions in an environment with a limited capacity for information acquisition.

Kempf, Manconi, and Spalt (2014) study the effect of experience at the industry level on manager performance. By working at the industry level, their work focuses on a broader concept of experience, while I highlight the expertise fund managers develop in specific companies. The magnitudes of my stock-level results are larger than their industry-level findings. Their clever identification relies on intra-portfolio differences, and credits a manager as gaining experience in an industry only when the manager holds a large portfolio share in an industry that suffers a period of poor industry level performance. This concept of accumulating experience only through a "trial by fire" – only by making mistakes, in some sense – is quite different from my own.

The impact of fund flows and large trades on asset prices has been robustly demonstrated. Coval and Stafford (2007) and Lou (2012) study this phenomenon in U.S. stock prices using mutual fund flows. Rhinesmith (2016) provides evidence of quarterly price impact for other trades that are large relative to volume. Notably, this effect would be expected to work against my findings. When a hedge fund establishes a large position in a stock at some point during a calendar quarter, which I utilize in my proxy for expertise, one would expect the price impact of those trades to increase the quarter end price of that stock. In other words, using only quarter end portfolio holdings and quarter end prices to measure risk-adjusted returns, one would expect fund flows to push towards finding manager underperformance, particularly on large holdings (since these holdings tend to have already borne larger investment flows by the time they show up in a 13F filing).

The impact of mutual fund flows on stock level volatility and comovement has been studied by Greenwood and Thesmar (2011). In part of this paper, I examine stock-level volatility with an eye towards available expert capital, rather than subscription and redemption-based mutual fund flows. I control for traditional measures of hedge fund and institutional ownership. However, my proxy for available expert capital may partially capture trading behavior of other investors that is not captured by my controls.

3.2. Data

I assemble data on hedge fund long U.S. stock holdings, stock returns, and company fundamentals.

I take hedge fund equity holdings from the Thompson Reuters Ownership Database (formerly the CDA/Spectrum Database), which compiles the 13F filings of all institutional investors. Any institution holding more than $100 million of 13F securities (mostly long U.S. equity holdings) must file a form 13F with the SEC each quarter.

I extract hedge fund 13F holdings from the full set of all institutional holdings using the comprehensive hedge fund dataset compiled by Agarwal, Fos, and Jiang (2013).[4] Since I am primarily examining the expertise of active, "stockpicking" managers, I use additional criteria to eliminate managers whose 13F filings are nonrepresentative of their investment strategies. For example, a fund that files only a single stock on a 13F report is likely primarily investing outside of publicly listed U.S.-equity holdings, while a fund that holds a controlling interest in a stock's common equity is likely pursuing a private equity strategy. Other funds, such as quant funds, are less likely to rely on stock-level expertise to make investment decisions and thus hold a relatively large number of positions. I therefore remove (1) any filing in which a single holding represents over 60% of the 13F portfolio, (2) any filing with fewer than 10 positions, (3) any filing in which a fund holds over 50% of the total outstanding market cap of a stock whose market cap exceeds $250 million, (4) any filing in which the value of the

[4] I thank Vikas Agarwal, Vyacheslav Fos, and Wei Jiang for providing me with this data.

13F portfolio is under \$50 million, and (5) any filing which contains more than 150 positions. None of my results are sensitive to these particular threshold values.

The 13F filings that remain do not capture complete fund portfolios. Notably, 13F filings do not provide information on short positions, cash holdings, or non-U.S. equity positions. However, after imposing the criteria above, the hedge fund 13F portfolios that remain should be representative of a manager's beliefs within that portfolio. If a manager holds a portfolio of 20 different stocks, in addition to cash, several short positions, and a number of credit positions, the manager's long stock positions should still represent, on average, a manager's beliefs within the stock portfolio. For instance, if a manager overweights a given stock within his long equity portfolio, it is reasonable to that the manager likely believes that that stock has some combination of greater expected return, lower risk, or some other characteristic that helps to attract or retain investors.

I obtain stock returns 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 obtain data on DGTW returns from Russ Wermers' website.[5] I eliminate the monthly returns of any stock that opens the month priced below \$5. I take the number of analysts covering a stock from I/B/E/S.

I examine 13F filings and stock return data from January 1, 1981 through June 30, 2012.

3.3. Expert positions

To construct an empirical proxy for a fund manager investing in a stock in which she is an "expert," I examine the event of a manager deciding to reenter the stock of a company that she already knows well. After reentry, I label such a manager-stock pair an "expert position."

I look for a sequence of three events to define an expert position. In chronological order,

[5]The DGTW benchmarks are available via http://www.smith.umd.edu/faculty/rwermers/ftpsite/Dgtw/coverpage.htm

first, a manager must establish expertise in a stock. I proxy expertise using position size as a percentage of total long equity holdings, under the assumption that a manager must develop expertise in a large position at some point before she sells the position. Second, the manager must effectively exit that stock, reducing the percentage portfolio position below some minimum size. At that point, the stock becomes "expert-eligible" for that manager. Finally, the manager must reenter that stock and once again make it a large position. Upon making this stock a large position, the manager-stock pair becomes an expert position.

The manager-stock pair continues to be an expert position as long as the position size exceeds the large position threshold. If the manager subsequently reduces the position size below the large position threshold, the manager-stock pair will cease to be an expert position, but will become an expert position again if it exceeds the large position threshold at any future time. In other words, after the manager has effectively exited the position for the first time, the stock is expert-eligible at all subsequent times. If the manager makes it a large position at any subsequent time, it will be an expert position.

Though I begin looking at 13F filings starting at January 1, 1981 to determine expert-eligibility, I only begin to examine expert positions starting on January 1, 1991 to allow "burn-in" for the procedure. Since a manager must first establish and then exit a position in a stock before the manager-stock pair is expert-eligible, there will be no expert positions at the beginning of my sample period (1981). Furthermore, the 13F hedge fund database is relatively small prior to 1991.

To distinguish a large position I define a size threshold at the manager level. In particular, the threshold I use is the greater of either 4% or 2.5 times a manager's average position size, defined as the inverse of the median number of positions on the manager's 13F filings.[6] I use an exit threshold of 0.25% to determine when a manager has effectively exited a position.

[6]In other words, if a manager files a median of 20 positions on her 13Fs, then her average position size is 1/20=5%. The large position threshold for such a manager would be 2.5*5%=12.5%. For a manager that files a median of 100 positions on her 13Fs, the threshold would be max(1/100*2.5=2.5%, 4%) = 4%.

As I illustrate in Section 3.5, my results are robust to varying these exact cutoffs.

After finishing this procedure, I have a set of 1,312,951 manager-stock-quarter observations, of which 61,138 represent large positions and 5,219 represent expert positions. Table 3.1 provides summary statistics.

Table 3.1. Hedge fund universe summary statistics

All statistics are taken across the set of 90 13F filings covered in the sample I use to form expert portfolios (12/31/1990 - 3/31/2012). Assets per quarter reflect the value of 13F U.S. equity holdings. 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.

	Mean	Median	Min	Max	Std Dev
Hedge funds per quarter	275.7	242	60	569	167.4
Positions per quarter	14499.8	14502	2877	27655	7427.9
Total assets per quarter ($ BB)	163.0	120.8	12.2	436.4	128.0
Median position value ($ MM)	2.2	2.1	1.1	3.3	0.5
Avg position size quintile*	3.8	3.9	3.5	4.2	0.2
Avg position book quintile*	2.7	2.7	2.5	3.1	0.1
Avg position momentum quintile*	3.1	3.1	2.8	3.5	0.2

3.3.1. Tendency to establish expert positions

I first examine the frequency with which fund managers establish expert positions. If expertise is valuable, and this expertise decays slowly, then we would expect managers to return to the stocks they know well more often than would a random placebo. Even if a manager knows a stock well, however, she could believe that it never becomes an attractive buy after she has exited it the first time. In other words, we would not expect a manager to establish an expert position in every expert-eligible stock.

Remarkably, of the set of all manager-stock pairs that managers make large positions at any point in time, managers establish expert positions in a full 8.9% of these manager-stock pairs. Equal-weighting across managers,[7] the average manager establishes expert positions

[7] In other words, I take the number of expert positions a manager enters during the sample, and divide that by

in 4.9% of the large positions that she initiates. The average manager initiates a total of 26 unique large positions over the sample, 2 of which become expert positions. In contrast, suppose that each time an average manager initiated a large position, or decided to establish an expert position, she had an equal probability of selecting any of the 2,000 most liquid stocks on the market.[8] In that case, we would expect a manager that initiates 28 total large positions (counting two initiations of the same stock as independent events) to establish 0.19 expert positions, or 0.7% of her total number of large positions. Such a manager would establish expert positions in 4.9% or more of her large positions only 1.4% of the time, giving a p-value of 0.014 for the observed value of the tendency of managers to establish expert positions. A more realistic approach, however, is to bootstrap the empirical distribution of the number of large positions that managers in the sample initiate. In other words, I assume that 40 managers initiate 5 large positions, 28 managers initiate 6 large positions, 29 managers initiate 7 large positions, etc., where these figures match the empirical distribution of the number of large positions that managers initiate in my sample. This produces a similar figure of 0.7% as the expected proportion of total large positions that managers make into expert positions. The p-value for seeing the average manager establish expert positions in 4.9% of her large positions, however, drops to 0.000 in this case.[9]

---

the number of large positions (positions that exceed the large position threshold) that the manager initiates over the sample. I then average across all managers. I exclude the 46 managers that never initiate a large position during the entire sample period.

[8]The average size quintile of all large positions held by managers in the sample is 4.1. This value makes it clear that there are numerous large positions selected from the second and third largest size quintiles. Using 2,000 to represent the size of the investment universe for large positions, which roughly corresponds to the two largest size quintiles, is conservative.

[9]This result is driven by the fact that there are a number of fund managers in the sample who initiate only a few large positions over the life of the fund. Since each time one of these managers initiates a large position there are fewer previous large positions to potentially select, these managers do not have as many opportunities to establish expert positions as a manager with many large positions. This effect is why equal weighting across managers produces a much lower figure for the percent of large positions that managers subsequently make expert positions. However, bootstrapping the empirical distribution of the number of large positions managers initiate takes this into account. Managers who only initiate a few large positions have virtually no probability of repeating these positions if they are choosing large positions "randomly." As a result, the p-value of the observed frequency of expert positions becomes effectively 0. This result holds even if one reduces the universe of potential large positions to a mere 500 stocks, an extremely conservative

These estimates are even more conservative because of the dynamic nature of establishing an expert position. A manager may have no or only a brief opportunity to make a large position initiated late in the fund's lifetime or near the end of the sample period into an expert position. For instance, at the extreme, if a manager initiates a large position and holds that position (i.e., never lets it fall below 0.25% of the long equity portfolio) until the fund's final 13F filing, the stock will never become expert-eligible. Following a similar line of reasoning, it is unlikely, albeit not impossible, that a manager will end up establishing an expert position in a stock that has only been expert eligible for a few quarters – one that the manager exited near the end of the fund's life or of the sample period – since this would mean that the manager would be reversing her sell decision quite quickly. This additional time dependence to determining expert-eligibility makes my treatment of the initiation of each large position as an independent event in my simulations a conservative assumption. Restricting the sample to managers that file at least 12 quarterly 13F reports, for instance, which reduces the sample to 600 managers and raises the ratio of the number of expert positions to the number of all large positions from 8.9% to 9.5%, or from 4.9% to 6.3% when equal-weighting across managers.

In summary, managers establish expert positions much more frequently than would be predicted by chance. I now turn to examining the performance of these expert positions.

3.3.2. Portfolio tests

3.3.2.1. Unconditional expert positions

I form calendar time portfolios to test whether managers outperform on expert positions.

I weight each expert position using its portfolio weight for the associated fund divided by the sum of the portfolio weights of all positions in the expert portfolio (so that final position

figure.

weights sum to 1).[10] For more details on this procedure, see Appendix A.1.

In Table 3.2 panel A, I display descriptive statistics of this portfolio across time. In Table 3.3 panel A, I display the risk- and characteristic-adjusted performance of the set of all hedge fund positions and of all large hedge fund positions in my sample.

In panel B, I show the performance of the expert portfolio. The portfolio exhibits positive but statistically insignificant outperformance. The four factor alpha of the portfolio is 2.98% on an annual basis, but its t-statistic is only 1.66. Thus, there is moderate evidence that fund managers know these specialist positions well. However, this effect is not statistically different from the 1.46% annualized four factor alpha of the portfolio of all stocks that funds hold as large positions, shown in panel A.

3.3.2.2. Separating the expert portfolio by trailing performance

However, fund managers may choose to establish expert positions for different reasons. On the one hand, a manager may believe that a stock is underpriced. On the other hand, the manager may have a dearth of investment ideas, but may want to appear to be maintaining an actively managed portfolio in order to justify active management fees. A natural way of differentiating between these two possibilities is to analyze the set of expert positions that are established after a recent decline in stock price. When a manager buys back into a stock after it has underperformed relative to the rest of the market, it may be more likely that the manager believes that the stock is undervalued. If the stock has been performing well recently, on the other hand, there may a greater chance it is a “feel good” stock that is easy to justify to or popular with investors, similar to window dressing in mutual funds.

I therefore form a portfolio of all expert positions for which, as of the quarter end at which the manager first established the stock as an expert position, the stock underperformed the

[10] I obtain similar results if I equal weight all expert positions.

Table 3.2. Expert portfolio summary statistics

All statistics are taken across the set of 90 13F filings covered in the sample I use to form expert portfolios (12/31/1990 - 3/31/2012). Assets per quarter reflect the value of 13F U.S. equity holdings. Quintile averages are equal weighted across all stocks. 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 A: Expert portfolio

	Mean	Median	Min	Max	Std Dev
Hedge funds per quarter	40.9	33	6	116	34.6
Positions per quarter	60.7	54	4	193	52.7
Total assets per quarter ($ BB)	5.25	3.87	0.16	22.66	5.5
Median position value ($ MM)	39.4	39.3	9.5	77.5	14.3
Avg position size quintile*	4.6	4.6	3.4	5.0	0.2
Avg position book quintile*	2.9	2.8	2.4	4.1	0.3
Avg position momentum quintile*	3.2	3.2	1.8	4.1	0.5

Panel B: Expert portfolio, after poor trailing performance

	Mean	Median	Min	Max	Std Dev
Hedge funds per quarter	20.4	12	2	67	18.3
Positions per quarter	24.3	18	2	80	22.1
Total assets per quarter ($ BB)	2.31	2.04	0.01	9.56	2.3
Median position value ($ MM)	31.2	30.6	1.7	86.0	16.3
Avg position size quintile*	4.5	4.5	3.3	5.0	0.3
Avg position book quintile*	2.9	2.9	1.7	4.7	0.5
Avg position momentum quintile*	2.9	2.8	1.5	5.0	0.7

Panel C: Expert portfolio, not after poor trailing performance

	Mean	Median	Min	Max	Std Dev
Hedge funds per quarter	27.4	24	1	86	23.7
Positions per quarter	36.2	34	1	116	31.2
Total assets per quarter ($ BB)	2.92	1.64	0.05	14.46	3.3
Median position value ($ MM)	46.7	44.6	9.8	96.3	17.6
Avg position size quintile*	4.7	4.7	3.4	5.0	0.2
Avg position book quintile*	2.9	2.8	2.1	4.5	0.5
Avg position momentum quintile*	3.5	3.5	2.0	5.0	0.7

Table 3.3. Aggregate hedge fund and expert portfolio performance

Monthly returns from 12/31/1990-6/30/2012 [t-statstics in brackets], based on the set of 90 13F filings covered in the sample I use to form expert portfolios (12/31/1990 - 3/31/2012).

Panel A: Hedge fund aggregated positions

	Raw return	DGTW adjusted	4 factor alpha	market	size	book	mom	CAPM alpha	market
all positions	1.01%	0.04%	0.06%	1.03	0.19	0.06	0.00	0.11%	1.06
		[1.47]	[1.29]	[90.70]	[13.00]	[4.11]	[-0.42]	[1.92]	[79.61]
large positions	1.10%	0.10%	0.12%	1.04	0.08	0.02	0.11	0.22%	1.02
		[1.59]	[1.40]	[50.15]	[3.03]	[0.59]	[6.60]	[2.39]	[48.97]

Panel B: Expert portfolio

Raw return	DGTW adjusted	4 factor alpha	market	size	book	mom	CAPM alpha	market
1.10%	0.12%	0.25%	0.98	-0.05	0.00	0.05	0.27%	0.95
	[0.97]	[1.66]	[27.55]	[-1.17]	[0.05]	[1.68]	[1.89]	[29.51]

median stock in my sample over the prior 6, 9, or 12 months.[11] For more detail on portfolio construction, see Appendix A.2. I similarly construct the returns of the portfolio of all expert positions that were not initiated after poor stock-level performance.

Panel B of Table 3.2 displays the summary statistics of the portfolio of expert positions initiated following stock-level underperformance. Table 3.4 panel A displays the performance results. Managers strikingly outperform on these stocks. The DGTW-adjusted performance, 4-factor alpha, and CAPM alphas of all of these portfolios are positive and economically and statistically significant. For example, the portfolio formed based on an expert position's trailing 12 month performance (prior to the manager establishing the expert position) generates DGTW-adjusted outperformance of 6.7%, an annualized four-factor alpha of 10.2%, or an annualized CAPM alpha of 8.3%. This outperformance is evidence that fund managers are able to identify attractive investment opportunities in expert stocks that have underperformed in the recent past.

Panel C of Table 3.2 displays the summary statistics of the portfolio of other expert positions. Table 3.4 panel B displays the performance results. Stocks that fund managers reinvest in after relatively strong past stock-level performance generate future risk-adjusted performance that is statistically indistinguishable from zero. DGTW-adjusted performance and four factor alphas are slightly negative, while CAPM alphas are slightly positive. However, if a manager is constrained by investors from passively indexing or holding too much cash, then the manager may still find it beneficial to invest in an expert stock that she does not expect to outperfom. For example, suppose that a manager first allocates capital to her true "alpha" opportunities, but that allocation is limited by price impact (Rhinesmith (2016)) or risk limits (Cohen, Polk, and Silli (2010)). She may therefore have a portion of her portfolio unallocated after exploiting her best opportunities. She needs to allocate this cap-

[11] When I measure underperformance using trailing 3 month return data, I find qualitatively similar results but lower point estimates and reduced statistical significance.

Table 3.4. Expert portfolios after splitting positions based on trailing performance

Monthly returns from 12/31/1990-6/30/2012 [t-statstics in brackets], based on the set of 90 13F filings covered in the sample I use to form expert portfolios (12/31/1990 - 3/31/2012).

Panel A: Expert positions established after stock-level performance over trailing interval < median stock return over that interval

Trailing return interval	Raw return	DGTW adjusted	4 factor alpha	market	size	book	mom	CAPM alpha	market
6 months	1.43%	0.46%	0.76%	0.96	-0.22	0.05	-0.21	0.59%	0.97
6 months		[2.37]	[3.07]	[16.13]	[-2.87]	[0.65]	[-4.28]	[2.32]	[16.94]
9 months	1.49%	0.54%	0.76%	0.95	-0.15	0.19	-0.20	0.66%	0.95
9 months		[2.62]	[3.05]	[15.96]	[-2.00]	[2.30]	[-4.14]	[2.54]	[16.42]
12 months	1.50%	0.54%	0.81%	0.92	-0.10	0.15	-0.24	0.67%	0.94
12 months		[2.41]	[2.98]	[14.14]	[-1.17]	[1.66]	[-4.49]	[2.42]	[15.16]

Panel B: Other expert positions

Trailing return interval	Raw return	DGTW adjusted	4 factor alpha	market	size	book	mom	CAPM alpha	market
6 months	0.95%	-0.05%	-0.06%	1.01	0.05	-0.02	0.24	0.12%	0.95
6 months		[-0.30]	[-0.34]	[24.53]	[0.97]	[-0.30]	[7.08]	[0.64]	[23.14]
9 months	0.89%	-0.10%	-0.06%	1.00	0.02	-0.11	0.23	0.06%	0.95
9 months		[-0.96]	[-0.37]	[23.60]	[0.33]	[-1.90]	[6.67]	[0.32]	[22.53]
12 months	0.91%	-0.09%	-0.05%	1.00	-0.03	-0.07	0.23	0.09%	0.94
12 months		[-0.67]	[-0.31]	[26.26]	[-0.57]	[-1.32]	[7.40]	[0.51]	[24.46]

ital to investments that appear "active" in order to justify her management fee to investors. Comparing a fairly priced expert stock and a random stock, both stocks will have similar expected risk-adjusted returns (approximately zero). However, given her prior knowledge of the company, the fund manager will be able to more effectively discuss the expert stock in meetings with investors. And demonstrating competence in such meetings is a critical step in attracting and preserving allocations from institutional investors.

3.3.2.3. Control portfolios

Could the strong outperformance generated by the former portfolio be based upon the simple strategy of buying stocks on a dip, rather than from the interaction of specialization with price dislocations? To test this concern, I examine the portfolio of all hedge fund positions that are initiated after recent stock-level underperformance in Table 3.5 panel A. See Appendix A.3 for further details on portfolio construction.

This portfolio generates small, generally insignificant positive values for risk-adjusted performance. Notably, the marginally positive risk-adjusted performance figures are very close to the positive risk-adjusted performance of the portfolio of all hedge fund positions, or of all large hedge fund positions, in my sample. The monthly four factor alpha is marginally significant, but at 0.09%-0.11% is economically much smaller than the outperformance of the expert portfolio of recently underperforming stocks. Neither the DGTW-adjusted performance or CAPM alpha are statistically distinguishable from 0.

As an additional control portfolio, I construct a characteristic-matched portfolio of potential reversal stocks. This portfolio represents the returns that a manager would reap by pursuing a reversal strategy with similar size and book-to-market characteristics as the portfolio formed in panel A of Table 3.4, but that is based solely on publicly available return data. I form this portfolio by first constructing the return series of 25 subportfolios. These portfolios are generated by a 5x5 dependent sort of all stocks by size and by indus-

Table 3.5. Control portfolios

Monthly returns from 12/31/1990-6/30/2012 [t-statstics in brackets], based on the set of 90 13F filings or set of stock returns covered in the sample I use to form expert portfolios. In panel A, the portfolio analyzes all positions that hedge funds in the sample purchase after the stock has had returns over the relevant trailing interval less than the median stock return over that period. In panel B, the portfolio is constructed as described in the text.

Panel A: All positions bought after poor trailing performance

Trailing return interval	Raw return	DGTW adjusted	4 factor alpha	market	size	book	mom	CAPM alpha	market
6 months	0.96%	0.03%	0.10%	1.02	0.13	0.08	-0.12	0.07%	1.07
6 months		[0.75]	[1.96]	[86.29]	[8.67]	[5.07]	[-11.92]	[1.00]	[72.18]
9 months	0.96%	0.03%	0.09%	1.02	0.14	0.10	-0.13	0.06%	1.07
9 months		[0.69]	[1.75]	[80.84]	[8.82]	[5.77]	[-12.44]	[0.83]	[66.19]
12 months	0.99%	0.05%	0.11%	1.02	0.16	0.12	-0.13	0.09%	1.07
12 months		[1.22]	[1.99]	[78.29]	[9.84]	[6.59]	[-11.91]	[1.16]	[63.48]

Panel B: Mechanical reversal strategy

Trailing return interval	Raw return	DGTW adjusted	4 factor alpha	market	size	book	mom	CAPM alpha	market
6 months	0.84%	-0.14%	0.07%	0.96	0.13	0.17	-0.25	-0.03%	1.03
6 months		[-1.49]	[1.24]	[75.11]	[7.96]	[9.58]	[-24.08]	[-0.30]	[44.04]
9 months	0.86%	-0.13%	0.06%	0.98	0.15	0.22	-0.28	-0.03%	1.04
9 months		[-1.19]	[0.99]	[66.15]	[7.87]	[11.15]	[-22.96]	[-0.25]	[39.00]
12 months	0.88%	-0.11%	0.07%	0.98	0.16	0.26	-0.27	0.00%	1.04
12 months		[-0.95]	[0.99]	[60.46]	[7.88]	[12.07]	[-20.59]	[-0.02]	[37.05]

try adjusted book-to-market, as in the construction of DGTW returns. For each of the 25 portfolios, I construct a reversal portfolio of the stocks in that size and value portfolio whose trailing 6, 9, or 12 month returns as of the end of any of the prior H quarter ends fell below the median stock return over the matched trailing return interval. This construction is equivalent to creating a portfolio consisting of all stocks with given size and book-to-market characteristics whose trailing returns are below average, and holding any stock that has been added to the portfolio for H quarters. I then market cap weight the stocks in each portfolio to generate return series for these 25 portfolios. See Appendix A.3 for further details on portfolio construction.

These alternative control portfolios also fail to generate outperformance, as shown for H=3 in Table 3.5 panel B.[12] These portfolios simply represent the returns to a strategy that is matched to my expert portfolio on size and book-to-market dimensions and that mechanically buys stocks after they fall. One can not replicate the strong returns of the expert portfolio using only information on past returns. Instead, in their expert eligible stocks fund managers appear to be able to distinguish a stock with an unwarranted price drop from a stock that will not recover in the future.

3.4. Available expert capital

If a fund manager accumulates a body of knowledge in the stocks in which she is an expert, then her decision to buy back into such a stock can carry a valuable endorsement. But what about the amount of capital that experts in a stock have available "on the sidelines"? If funds actively monitor their expert stocks, then if a mispricing forms in such a stock, those funds may trade to exploit that mispricing. In turn, other market participants may be aware

[12] These results are robust to other alternative holding periods, which I do not display to conserve space. Using holdings periods of 2 or 4 quarters (H=2 or H=4), the characteristic-matched reversal portfolio also fails to outperform.

of the presence and capital of experts. The amount of expert capital that is available to be deployed into a given stock may thus impact asset price behavior.

Empirically, I find that available expert capital forecasts stock-level volatility. In Section 3.4.1, I explain the empirical approach and results. In Section 3.4.2, I suggest possible interpretations of this finding.

3.4.1. Forecasting stock-level volatility

I define available expert capital by summing the total 13F assets of all funds for whom a given stock is expert-eligible. This figure represents the amount of assets that funds could theoretically deploy into a given expert-eligible position. See Appendix B for details.

Figure 1 presents return volatilities – measured as the standard deviation of daily or 10-day returns over the quarter following the date at which I measure available expert capital – of 10 portfolios sorted into available expert capital deciles. Since a large number of stocks have zero available expert capital, I omit them from this figure. Clearly, stocks with higher available expert capital exhibit greater return volatility.

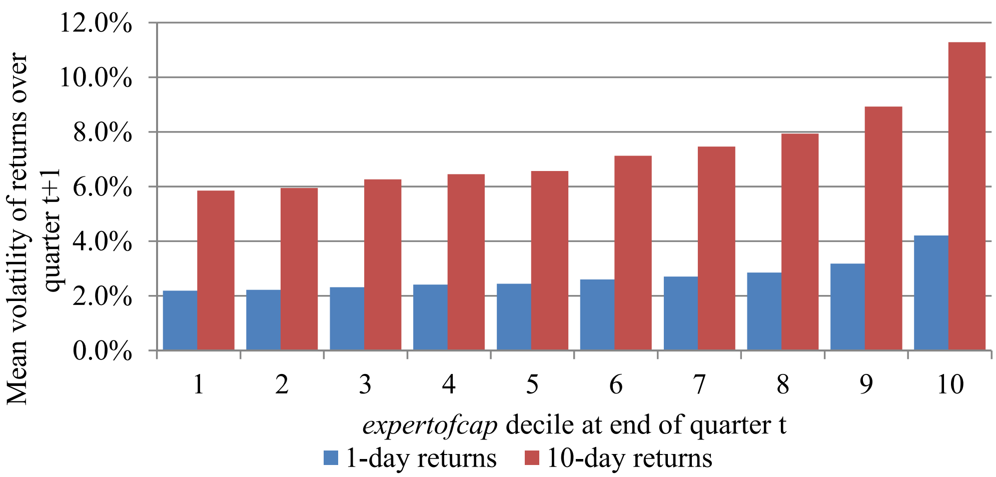

Figure 3.1. Stock return volatility and available expert capital deciles

I employ two specifications to formally test whether available expert capital predicts return volatility. In the first, I use a Fama-Macbeth regression approach.

$$\sigma_{s,t+1} = \alpha + \beta available\, expert\, capital_{s,t} + \gamma_1 ownershipcontrols_{s,t} + \gamma_2 othercontrols_{i,t} + \epsilon_{s,t}$$

In the second, I employ firm-level fixed effects with standard errors clustered by time period.

$$\sigma_{s,t+1} = \delta_s + \beta available\, expert\, capital_{s,t} + \gamma_1 ownershipcontrols_{s,t} + \epsilon_{s,t}$$

Table 3.6 presents the results.[13] The coefficient on available expert capital is positive and statistically significant under all specifications. Using stock fixed effects in the sample of only observations with non-zero available expert capital, I find that a one standard devation increase in available expert capital is associated with an increase in next quarter's volatility of 1-day returns of 0.32%, which is 21.6% of the standard deviation or 15.3% of the mean of the dependent variable, the standard deviation of a stock's 1-day returns over the following quarter. Fama-Macbeth estimates are lower in magnitude, but remain highly significant.

Notably, these regressions control for both the level and concentration of overall hedge fund ownership in a stock (in particular, the log of the number of hedge fund owners, the percentage of market cap owned by hedge funds, and the Herfindahl index of hedge fund ownership stakes), as well as the level of general institutional ownership (the percentage of market cap owned by all 13F filers), which means that the available expert capital coefficient does not simply reflect a stock's contemporaneous ownership base. The Fama-Macbeth approach additionally controls for, at the stock-quarter level, the log of market capitalization, the log of the number of sell-side analysts, book-to-market, the log of price, volume, the log of age, trailing 12 month returns, trailing 12 month skewness, and lagged volatility.

These results exploit variation in available expert capital across stocks and/or time. If hedge funds simply prefer more volatile stocks, or if hedge funds try to predict volatility and thus buy stocks before volatility increases, my ownership controls would pick that up.

[13] I winsorize available expert capital, volatility, and fund flow measures at 99.5%.

Table 3.6. Available expert capital regressions

Regressions as described in the text, using volatility of returns returns from 12/31/1990-6/30/2012 [t-statstics in brackets]. Available expert capial is based on the set of 90 13F filings covered in the sample I use to form expert portfolios (12/31/1990 - 3/31/2012).

	Volatility of total returns in the following quarter, with returns measured over a horizon of:							
Dependent variable	10 days	10 days	10 days	10 days	1 day	1 day	1 day	1 day
Inc. available expert capital=0	N	N	Y	Y	N	N	Y	Y
Available expert capital	0.0020	0.0080	0.0013	0.0050	0.0005	0.0032	0.0002	0.0023
	[6.13]	[6.33]	[5.08]	[3.81]	[6.38]	[5.86]	[3.23]	[4.32]
num. of hedge	0.0002	-0.0002	0.0003	-0.0003	0.0000	0.0000	0.0001	-0.0001
fund owners	[5.34]	[-4.42]	[8.91]	[-4.30]	[5.73]	[-2.77]	[9.87]	[-2.64]
hedge fund	-0.0002	-0.0370	-0.0134	-0.0276	-0.0020	-0.0148	-0.0085	-0.0106
share mkt cap	[-0.02]	[-3.88]	[-4.27]	[-2.91]	[-0.95]	[-4.33]	[-9.17]	[-2.83]
inst. share	-0.0044	-0.0371	0.0020	-0.0339	-0.0002	-0.0120	0.0015	-0.0156
mkt cap	[-2.84]	[-3.62]	[1.63]	[-5.76]	[-0.45]	[-2.79]	[4.58]	[-7.07]
hedge fund	0.0005	0.0064	0.0003	0.0013	0.0002	0.0018	-0.0002	-0.0011
Herfindahl Index	[0.05]	[3.02]	[0.95]	[1.78]	[1.06]	[2.27]	[-2.58]	[-3.69]
constant	0.0514	0.0947	0.0861	0.1117	0.0157	0.0335	0.0311	0.0476
	[10.99]	[12.51]	[29.23]	[36.08]	[12.91]	[10.85]	[33.54]	[48.87]
Additional controls	Y	N	Y	N	Y	N	Y	N
R^2	0.38	0.36	0.36	0.35	0.56	0.40	0.59	0.49
estimation	FM	Fixed effects	FM	Fixed effects	FM	Fixed effects	FM	Fixed effects

Instead, the coefficient on available expert capital relies on the variation that occurs if, for example, an expert fund that does not currently own stock X performs well on the rest of its portfolio or has large investment inflows.

In sum, the available amount of expert capital in a stock robustly predicts the future realized volatility of returns in that stock.

3.4.2. Interpretation: more adverse selection or predicting flow-driven price changes?

I suggest two potential interpretations of the result that available expert capital predicts future return volatility. In the first interpretation, available expert capital increases the probability of informed trading in a stock. As a result, market makers increase spreads, causing larger bid-ask bounces and greater realized volatility. In the second interpretation, available expert capital reflects the increased vulnerability of certain stocks to correlated shifts in investor demand. The resulting trades drive non-fundamental price volatility in stocks.

The reasoning for the first channel is that if expert capital is informed, then with more expert capital one would expect a higher proportion of information-based trading. Intuitively, market makers would then demand larger spreads to compensate for the possibility of adverse selection, as in Glosten annd Milgrom (1985). As a result, even uninformed trades will now induce greater return volatility as they push stock prices across wider spreads. This phenomenon translates into higher realized return volatility over quarter t+1 in stocks in which there is more expert capital available at the end of quarter t.[14]

[14]In this interpretation, one might also wonder about the ability of available expert capital to predict structural measures of the amount of asymmetric information in a stock such as PIN (Easley, Kiefer, O'Hara, and Paperman 1996) or adjusted PIN (Duarte and Young 2009). Indeed, in unreported results, available expert capital is able to positively and significantly predict PIN, and is able to positively predict adjusted PIN, although this latter finding is significant in only some specifications. However, I emphasize return volatility because it represents an intuitive and non-parametric measure of asset price dynamics. The structural

The second channel appeals to the underlying correlated network of the fund flows and trades of institutional investors, as in Greenwood and Thesmar (2011). In this setting, available expert capital may predict large future imbalances in investor demand unrelated to fundamentals. Perhaps available expert capital is associated with stocks that are more susceptible to changes in sentiment by uninformed mutual funds, for example. In a setting with downward sloping demand curves, this could lead to elevated return volatility.

In future work, I intend to consider empirical approaches that may distinguish between these two channels. For example, perhaps long-horizon return dynamics or price reactions to earnings announcements (as in Rhinesmith (2016)) may help separate information-based from non-fundamental price movements. Similarly, examining future total trading volumes and flow-driven directional trades may be relevant.

3.5. Robustness checks

The risk-adjusted outperformance of expert positions that are bought after poor trailing stock-level performance is not sensitive to the cutoff values used. It is also stable over time, albeit weaker during the second half of the sample.

3.5.1. Cutoffs

In Table 3.7, I display the sensitivity of my portfolio returns to the various cutoff values employed. I vary both the minimum position size on both an absolute (% of portfolio, panel A) and relative (multiple of that manager's average position size, panel B) basis. I display

models underlying PIN measures, in contrast, assume rigidities in the trading process that I do not believe are appropriate here. First, PIN models do not allow for non-information based liquidity events to drive consistent order imbalances. For example, the literature has demonstrated that redemption and subscription driven mutual fund flows drive order flow without reflecting private information. Second, PIN models assume that daily trading volumes and order imbalances will allow one to differentiate between informed and uninformed trading. However, stock experts may trade across horizons longer than a day given their tendency to build large positions (see, for example, Rhinesmith (2016) or Collin-Dufresne and Fos (2015)).

only measures of risk adjusted performance, and their associated t-statistics, for convenience. As one can see, the portfolio of expert positions bought on a dip still generates statistically and economically significant outperformance across nearly all specifications. In several cases, the point estimates are higher than in my baseline setup.

3.5.2. Subsample stability

The risk adjusted outperformance of the expert portfolio formed of expert positions initiated afer stock-level underperformance remains across sample periods, albeit it is weaker (but more precisely estimated) in the second half of the sample. Splitting the sample into observations in the 1990s and observations after January 1, 2000, for instance, I get annualized four-factor alphas of 14.9% and 5.8%, with t-statistics of 2.18 and 1.97, respectively, using a six-month lookback period.

3.6. Conclusion

Hedge fund managers display an outsized tendency to repurchase stocks which they have previously held as large positions – managers appear to specialize in particular stocks. When a fund purchases a stock in which it is an "expert" after that stock has recently underperformed, it is a credible signal that the stock is undervalued. A portfolio formed on such a basis has historically generated annualized risk-adjusted outperformance of between 5% and 10%. Furthermore, variation in the available expert capital tied to a given stock reliably predicts an increase in the volatility of the future returns of that stock. This result may reflect an adjustment in the price process to the presence of more informed traders, or it may be that expert capital predicts correlated shifts in uninformed demand. Future empirical work may help distinguish between these channels.

My findings on stock experts underline the importance of investment relationships in

Table 3.7. Alternative expert portfolios after poor trailing performance

Monthly returns from 12/31/1990-6/30/2012 [t-statstics in brackets], based on the set of 90 13F filings covered in the sample I use to form expert portfolios (12/31/1990 - 3/31/2012).

Panel A: Expert portfolio returns, changing % cutoff

Trailing return interval	DGTW adjusted			4 factor alpha			CAPM alpha		
	6 mths	9 mths	12 mths	6 mths	9 mths	12 mths	6 mths	9 mths	12 mths
3%				0.65%	0.54%	0.56%	0.52%	0.45%	0.45%
cutoff	[2.21]	[2.34]	[2.58]	[3.10]	[2.69]	[2.91]	[2.35]	[2.13]	[2.19]
3.5%	0.41%	0.44%	0.50%	0.70%	0.65%	0.70%	0.57%	0.56%	0.58%
cutoff	[2.31]	[2.62]	[2.80]	[3.14]	[3.14]	[3.27]	[2.47]	[2.57]	[2.58]
4.5%	0.53%	0.49%	0.49%	0.82%	0.78%	0.74%	0.64%	0.67%	0.61%
cutoff	[2.35]	[2.08]	[1.94]	[2.99]	[2.73]	[2.49]	[2.29]	[2.28]	[2.00]
5%	0.49%	0.50%	0.60%	0.74%	0.81%	0.80%	0.55%	0.67%	0.67%
cutoff	[1.69]	[1.57]	[1.93]	[2.27]	[2.26]	[2.35]	[1.67]	[1.84]	[1.94]

Panel B: Expert portfolio returns, changing relative cutoff

Trailing return interval	DGTW adjusted			4 factor alpha			CAPM alpha		
	6 mths	9 mths	12 mths	6 mths	9 mths	12 mths	6 mths	9 mths	12 mths
2x avg	0.31%	0.32%	0.30%	0.65%	0.58%	0.61%	0.59%	0.49%	0.55%
position	[1.71]	[1.75]	[1.45]	[2.94]	[2.65]	[2.55]	[2.04]	[1.76]	[1.73]
2.25x avg	0.34%	0.41%	0.41%	0.69%	0.68%	0.71%	0.60%	0.52%	0.60%
position	[1.75]	[2.07]	[1.91]	[2.89]	[2.87]	[2.85]	[1.96]	[1.81]	[1.82]
2.75x avg	0.55%	0.64%	0.63%	0.75%	0.79%	0.84%	0.53%	0.62%	0.65%
position	[2.61]	[2.91]	[2.55]	[2.84]	[2.99]	[2.79]	[1.64]	[2.16]	[1.94]
3x avg	0.40%	0.46%	0.49%	0.64%	0.65%	0.70%	0.56%	0.42%	0.53%
position	[1.64]	[1.82]	[1.66]	[2.31]	[2.33]	[2.14]	[1.84]	[1.46]	[1.52]

capital formation. If the network of relationships matters, then both the level and the distribution of capital – or of capital losses during crisis episodes – across investment funds may impact outcomes.

Bibliography

[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] Alexander, Gordon, Gjergji Cici, and Scott Gibson, 2007, Does motivation matter when assessing trade performance? An analysis of mutual funds, Review of Financial Studies 20, 125-150.

[4] Amihud, Yakov, 2002, Illiquidity and stock returns: Cross-section and time-series effects, Journal of Financial Markets 5, 31–56.

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

[6] 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.

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

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

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

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

[11] 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.

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

[13] 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.

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

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

[16] 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.

[17] 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.

[18] 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.

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

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

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

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

[23] 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.

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

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

[26] Duarte, Jefferson and Lance Young, 2009, Why is PIN priced?, Journal of Financial Economics 91, 119-138.

[27] Easley, David, Nicholas Kiefer, Maureen O'Hara, and Joseph Paperman, 1996, Liquidity, information, and infrequently traded stocks, Journal of Finance 51, 1405–1436.

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

[29] Frazzini, Andrea, 2006, The disposition effect and underreaction to news, Journal of Finance 61, 2017-2046.

[30] 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.

[31] Gamboa-Cavazos, Mario and Pavel Savor, 2011, Holding on to your shorts: When do short sellers retreat?, Working paper.

[32] Greenwood, Robin and David Thesmar, 2011, Stock price fragility, Journal of Financial Economics 102, 471-490.

[33] 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.

[34] 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.

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

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

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

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

[39] 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.

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

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

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

[43] Kempf, Elisabeth, Alberto Manconi and Oliver Spalt, 2014, Learning by doing: The value of experience and the origins of skill for mutual fund managers, Working paper.

[44] 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.

[45] 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.

[46] 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.

[47] Kruttli, Mathias, Andrew Patton and Tarun Ramadorai, 2014, The impact of hedge funds on asset markets, Working paper.

[48] Koudijs, Peter, 2014, Those who know most: Insider trading in 18th c. Amsterdam, Working paper.

[49] Kumar, Alok and Charles Lee, 2006, Retail investor sentiment and return comovements, Journal of Finance 61, 2451-2486.

[50] Kyle, Albert, 1985, Continuous auctions and insider trading, Econometrica 53, 1315-1335.

[51] 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.

[52] Lou, Dong, 2012, A flow-based explanation for return predictability, Review of Financial Studies 25, 3457-3489.

[53] Di Mascio, Rick, Anton Lines and Narayan Naik, 2015, Alpha decay, Working paper.

[54] Merton, Robert, 1987, A Simple Model of Capital Market Equilibrium with Incomplete Information, Journal of Finance 42, 483–510.

[55] Pastor, Lubos and Robert Stambaugh, 2012, On the size of the active management industry, Journal of Political Economy 120, 740-781.

[56] Pastor, Lubos, Robert Stambaugh, and Lucian Taylor, 2015a, Scale and skill in active management, Journal of Financial Economics 116, 23-45.

[57] Pastor, Lubos, Robert Stambaugh, and Lucian Taylor, 2015b, Do funds make more when they trade more?, Working paper.

[58] Patton, Andrew, Tarun Ramadorai, and Michael Streatfield, 2015, Change you can believe in? Hedge fund data revisions, Journal of Finance, 70, 963-999.

[59] Pollet, Joshua and Mungo Wilson, 2008, How does size affect mutual fund behavior?, Journal of Finance 63, 2941-2969.

[60] Rhinesmith, Jonathan, 2014, Doubling Down, Working paper.

[61] Scharsftein, David and Jeremy Stein, 1990, Herd behavior and investment, American Economic Review 80, 465-479.

[62] Shleifer, Andrei and Robert Vishny, 1997, The limits of arbitrage, Journal of Finance 52, 35-55.

[63] Van Nieuwerburgh, Stijn, and Laura Veldkamp, 2010, Information acquisition and under-diversification, Review of Economic Studies, 77, 779-805.

[64] Wermers, Russ, 2004, Is money really "smart"? New evidence on the relation between mutual fund flows, manager behavior, and performance persistence, Working paper.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.

[65] Wurgler, Jeffrey and Ekaterina Zhuravskaya, 2002, Does arbitrage flatten demand curves for stocks?, Journal of Business 75, 583-608.

A. Appendix to Chapter 1

A.1. Proofs

A.1.1. Additional model equations

The following equilibrium equations constrain the parameters of the model:

$$
\begin{aligned}
\beta_1 &= \frac{1}{2\lambda_1}\left(\frac{\lambda_2 - \frac{1}{2}\lambda_1\pi}{\lambda_2 - \frac{1}{4}\lambda_1\pi}\right) && \text{(A.1)} \\
\beta_2 &= \frac{1}{2\lambda_2} && \text{(A.2)} \\
\lambda_1 &= \frac{\beta_1\phi\sigma_\epsilon^2}{\beta_1^2\phi^2(\sigma_\epsilon^2 + \sigma_\eta^2) + \sigma_u^2} && \text{(A.3)} \\
\lambda_2 &= \frac{\beta_2\phi(1 - \lambda_1\beta_1\phi)\sigma_\epsilon^2}{\beta_2^2\phi^2(1 - \lambda_1\beta_1\phi)(\sigma_\epsilon^2 + \sigma_\eta^2) + \sigma_u^2} && \text{(A.4)} \\
\phi &= \frac{\sigma_\epsilon^2}{\sigma_\epsilon^2 + \sigma_\eta^2} && \text{(A.5)}
\end{aligned}
$$

A.1.2. Model solution

To solve the model, I conjecture equations (1.1)-(1.4) and (A.1)-(A.5). 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 (1.2), with $\beta_2 = \frac{1}{2\lambda_1}$ as in (A.2). 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 (1.1), with β_1 as in (A.1). 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 (1.3), with λ_1 as in (A.3).

   b) Solve for p_2 also using Bayesian updating and assuming that y_1 and y_2 are independent. The result is (1.4), with λ_2 as in (A.4). Note further that (A.4) 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.1.3. 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.1.4. 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 $\frac{1}{2\beta_1}(1-\beta_1\lambda_1)x_1$.

$\beta_1\lambda_1 \le \frac{1}{2}$ 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.1.5. 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.1.6. 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.1.7. 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.1.8. 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 (A.3) to eliminate all λ_1 terms. Obtain the equation $a_3*(2\sqrt{a_1a_2}-\pi\sigma_u\sqrt{a_3}) = a_2*(\sqrt{a_1a_2}-\pi\sigma_u\sqrt{a_3})$ 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 (A.2) and (A.4). Thus $E[mispricing] = constant * \frac{x_2}{\sigma_u}$.

A.2. Example and robustness to alternative explanations

A.2.1. 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.[1] 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.[2] 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 large trades predict future returns. Furthermore, I present evidence that the price impact of these trades incorporates information into prices as the trades occur.

A.2.2. Alternative explanations

Volume consumed predicts future performance beyond existing alternative explanations. Table A.1 shows future performance (quarter t+1) after controlling for each alternative. To

---

[1] 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.

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

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 1.2 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 A.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 1.7.2, 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.[3]

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

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

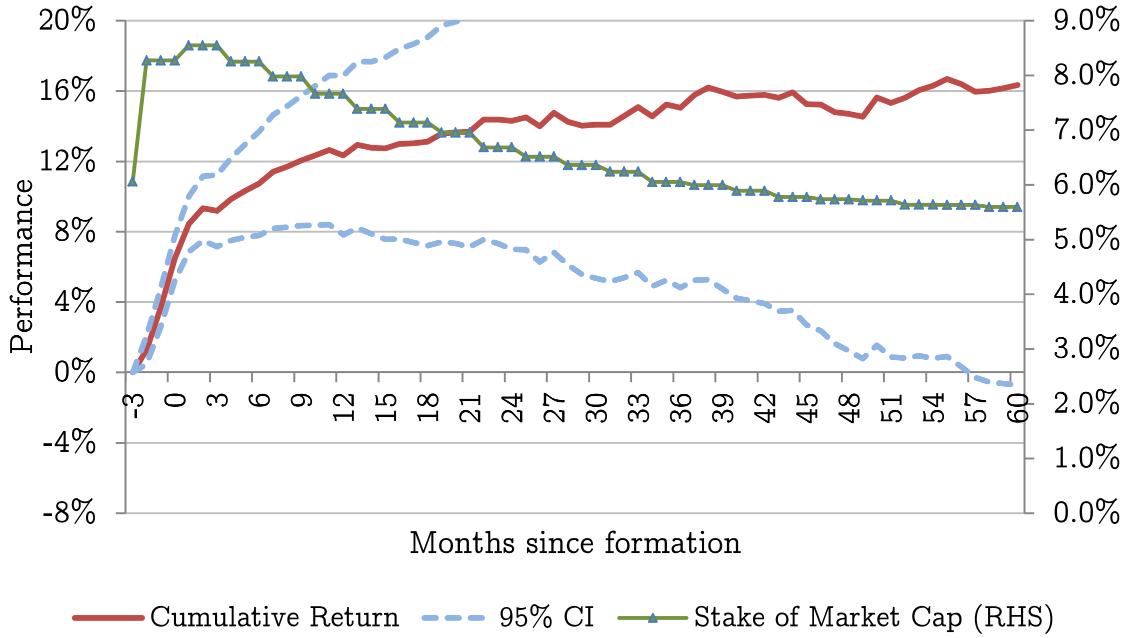

Figure A.1. Volume consumed, cumulative performance and sales over time
This figure displays the cumulative buy and hold performance of portfolios that go long stocks in the top decile of hedge fund volume consumed (aggregation method 1) and short stocks in the lowest decile. It also displays the average 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.

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 A.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.

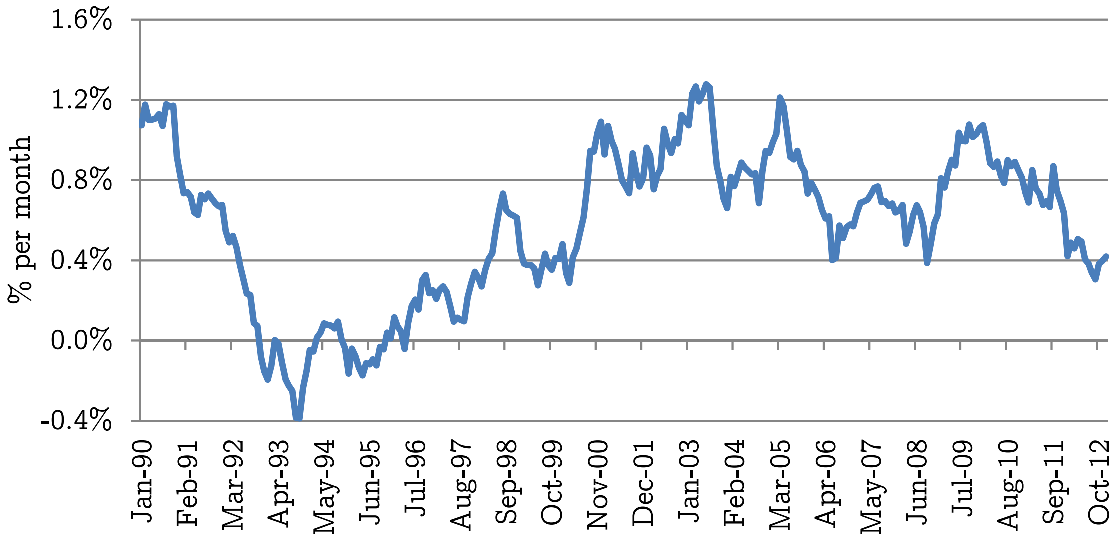

Figure A.2. Long-short portfolio, 3-year trailing performance

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 A.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 A.2.2. 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.

Table A.1: (continued)

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 / asymmetric info		
No top decile					
of managers by avg.	0.49%	0.46%	Double sort:	0.63%	0.46%
stake in company	[3.71]**	[5.31]**	volat 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 >	0.41%	0.32%
future 1 quarter flows	[4.00]**	[4.78]**	median NYSE	[3.45]**	[4.15]**
Normalize by mkt cap					
Double sort:	0.33%	0.26%	Only $ volume >	0.74%	0.59%
Market cap consumed	[1.78]*	[4.10]**	median NYSE	[5.86]**	[6.28]**

A.3. Maximum likelihood estimation

A.3.1. Building the likelihoood function

Let $\Theta = (\Theta_1; \Theta_2) = (\sigma_\epsilon^2, \sigma_\eta^2, \sigma_u^2, \pi;\ \beta_1, \lambda_1, \beta_2, \lambda_2, \phi)$ be the vector of parameters. Let $X = (x_{1,s,t}, x_{2,s,t}, r_{1,s,t}, r_{2,s,t})$ be the vector of observables, with $r_2 = p_2 - p_1$ and $r_1 = p_1$. Let $\mathbf{1}(x)$ be an indicator variable equal to 1 if the event x has occured, and 0 otherwise. Let $g(x)$ be the standard normal PDF and $G(x)$ the standard normal CDF.

Solve for the probability of X given Θ. s indexes stocks, t indexes information episodes (rather than quarters), and the subscript of 1 (subscript of 2) denotes the first (second) quarter in each episode. The likelihood function for observing X given Θ, with x_1 and x_2 censored below at x_c and information publicly released after quarter 1 (as opposed to after quarter 2) if $x_1 > x_c$ and $x_2 \leq x_c$, is:

$$likelihood(X|\theta) = \prod_{t=1}^{T}\prod_{s=1}^{S} \mathbf{1}(x_{1,s,t} > x_c\, and\, x_{2,s,t} \leq x_c) * Pr(X_1 = x_{1,s,t}\, and\, R_1 = r_{1,s,t}) +$$

$$\mathbf{1}(x_{1,s,t} > x_c\, and\, x_{2,s,t} > x_c) * Pr(X_1 = x_{1,s,t}\, and\, R_1 = r_{1,s,t}\, and\, X_2 = x_{2,s,t}\, and\, R_2 = r_{2,s,t})$$

$$+ \mathbf{1}(x_{1,s,t} \leq x_c) * Pr(X_1 \leq x_c) \quad (A.6)$$

Now solve and plug in for the vector of unobserved random variables $(i_{s,t}, u_{1,s,t}, u_{2,s,t})$ as a function of X and Θ. Combined with taking the log of the likelihood function, this produces:

$$log(likelihood(X|\Theta)) = \sum_{t=1}^{T}\sum_{s=1}^{S} \mathbf{1}(x_{1,s,t} > x_c\, and\, x_{2,s,t} \leq x_c) * \Big\{ log(g(\frac{r_{1,s,t} - \lambda_1 x_{1,s,t}}{\lambda_1 \sigma_u}))$$

$$+ log(g(\frac{x_{1,s,t}}{\beta_1 \phi \sqrt{\sigma_\epsilon^2 + \sigma_\eta^2}}))\Big\} + \mathbf{1}(x_{1,s,t} > x_c\, and\, x_{2,s,t} > x_c) * \Big\{ log(g(\frac{r_{1,s,t} - \lambda_1 x_{1,s,t}}{\lambda_1 \sigma_u})) + log(g(\frac{r_{2,s,t} - \lambda_2 x_{2,s,t}}{\lambda_2 \sigma_u}))$$

$$+ log(g(\frac{x_{1,s,t}}{2\beta_1 \phi \sqrt{\sigma_\epsilon^2 + \sigma_\eta^2}} + \frac{x_{2,s,t} + \beta_2 r_{1,s,t}}{2\beta_2 \phi \sqrt{\sigma_\epsilon^2 + \sigma_\eta^2}}))\Big\} + \mathbf{1}(x_{1,s,t} \leq x_c) * G(\frac{x_c}{\beta_1 \phi \sqrt{\sigma_\epsilon^2 + \sigma_\eta^2}}) \quad (A.7)$$

The model is also subject to constraints, equations (A.1)-(A.5), which implicitly define Θ_2 by $f_1(\Theta_1;\ \Theta_2) = 0$. Equation (A.5) is trivial. Equations (A.1)-(A.4) can be reduced to single implicit equation for $f_2(\Theta_1; \beta_1) = 0$:

$$a_3 * (2\sqrt{a_1 a_2} - \pi\sigma_u\sqrt{a_3}) = a_2 * (\sqrt{a_1 a_2} - \pi\sigma_u\sqrt{a_3}) \tag{A.8}$$

with $a_1 = \sigma_u^2 + \beta_1^2\phi^2\sigma_\eta^2$, $a_2 = \sigma_u^2 + \beta_1^2\phi\sigma_\epsilon^2$, and $a_3 = \beta_1^2\phi\sigma_\epsilon^2$. Thus given Θ_1, one can numerically solve for β_1. Once one has solved for β_1, one can explicitly obtain $\lambda_1 = f_3(\beta_1)$, $\lambda_2 = f_4(\beta_1, \lambda_1)$, and $\beta_2 = f_5(\lambda_2)$ in turn.

A maximum likelihood approach is implemented by maximizing equation (A.7) over Θ_1, subject to the constraints on Θ_2 (equation (A.8)), given data X.

I use maximum likelihood because it allows me to both utilize the Kyle model's explicit structure on error terms (i.e., u_1 and u_2 are normal and the market maker infers information from order flow based on that parametrization) and to model censoring (by integrating those errors terms over the relevant range for censored data). In the likelihood function, in the case that both $x_{1,s,t}$ and $x_{2,s,t}$ are observed above the point of censoring I combine information from equations (1.1) and (1.2) to solve for $i_{s,t}$. In order do so, I utilize a modeling technique to address the fact that the model imposes a restriction on the *data* – not the parameters – that does not hold exactly. The issue is that both equations (1.1) and (1.2) can be solved for ϕi. To proceed, I rewrite (1.1) and (1.2) with a noise term added to each, $\xi_{1,s,t}$ and $\xi_{2,s,t}$, respectively. I assume both of these variables are i.i.d. $N(0, \sigma_\xi^2)$ and independent of all other random variables. I take $\sigma_\xi^2 \ll \beta_i^2(\sigma_\epsilon^2 + \sigma_\eta^2)$, $i = 1, 2$, so that the likelihood function ignores this noise (since it is so small relative to other sources of variation). The intuition I hope to gain from the structural model is based on *imposing* that hedge fund trades are based on information. Rather than introducing a free parameter for noise in hedge fund trading, I introduce a minimal amount of noise to informed trading to implement the model. The result of this technique is that if both $x_{1,s,t}$ and $x_{2,s,t}$ are observed above the point of censoring, the model simply averages the information they contain for $i_{s,t}$ (from equations

(1.1) and (1.2)).

A.3.2. CARA informed trader and public information shocks

I solve the model with risk aversion and new information shocks.

First, I assume the informed trader has CARA utility. Second, I assume that public "new information" arrives during each period that is independent of the original information draw. This information is public in the sense that it moves prices *without* trading as soon as it is generated. Note that a risk neutral trader would simply ignore such information (if it is mean zero) in her optimization.

Specifically, the informed trader has utility $U = -e^{-aW}$, where W is her wealth at the end of the current two-quarter information event. The informed trader's maximization now must take into account the risk generated by noise trading and the public new information events. The informed trader maximizes $E[profits] - \frac{a}{2}\, variance(profits)$.

Denote the public new information shock each period as $ni_t \sim N(0, \sigma_{ni}^2)$. Profits from the second period of trading are $x_2(\epsilon - p_1 - \lambda_2 u_2 + ni_2) - x_2^2\lambda_2$. Trading during period 2 is assumed to take place before the price shock ni_2, and hence new information in period 2 affects profits on the quantity traded during period 2 (x_2). Choosing x_2 to maximize the mean minus ($\frac{a}{2}*$) the variance of this quantity produces $x_2 = \beta_2(\phi i - p_1)$, with $\beta_2 = 1/D$, and $D = \{2\lambda_2 + a\lambda_2^2\sigma_u^2 + a\sigma_{ni}^2 + a(1-\phi)\sigma_\epsilon^2\}$. Note from this equation that, strictly speaking, with risk aversion the optimal amount of informed trading is no longer linearly related to the mispricing divided by the expected magnitude of noise trading. The former enters linearly, but the latter does not. As a result, my empirical proxy for x_t may not be as effective in this extension.

Stepping back to the first period, the informed trader maximizes the mean minus ($\frac{a}{2}*$) the variance of $x_1(\epsilon - \lambda_1 u_1 + ni_2 + ni_1) - x_1^2\lambda_1 + \pi * (profits\, from\, second period)$ over x_1. The quantity traded during period 1 is assumed to be subject to both price shocks ni_1 and

ni_2.

This produces $x_1 = \beta_1(\phi i)$, with β_1 satisfying the following equation:

$$\beta_1 * \{-(-2\lambda_1 D^4 + 2\lambda_1^2 \pi D^3 - 2\pi\lambda_1^2\lambda_2 D^2) + a\{(1-\phi)\sigma_\epsilon^2(D^2 - \pi\lambda_1 D)^2 +$$
$$\sigma_u^2(\pi^2\lambda_1^2\lambda_2^2 D^2 + (-\lambda_1 D^2 + 2\pi\lambda_1^2 D - 2\pi\lambda_1^2\lambda_2)^2 + \sigma_{ni}^2(D^2 - \pi\lambda_1 D)^2 + \sigma_{ni}^2 D^4\}\}$$
$$= \{D^4 - 2\pi\lambda_1 D^3 + 2\pi\lambda_1\lambda_2 D^2 - a\{(1-\phi)\sigma_\epsilon^2\pi(D^3 - \pi\lambda_1 D^2) +$$
$$\sigma_u^2(-\pi^2\lambda_1\lambda_2^2 D^2 + (-\lambda_1 D^2 + 2\pi\lambda_1^2 D - 2\pi\lambda_1^2\lambda_2)(-\pi\lambda_1 D + 2\pi\lambda_1\lambda_2)) + \sigma_{ni}^2\pi(D^3 - \pi\lambda_1 D^2)\}\} \tag{A.9}$$

The market maker proceeds as before, given that the informed trader will trade an amount proportional to β_t times the remaining mispricing.

The likelihood function is the same as above. The constraints, however, can no longer be reduced to a single constraint. Instead, I numerically solve the revised constraint for β_1, (A.9), jointly with equations (A.3) and (A.4). Given β_1, λ_1, and λ_2, I calculate D, which in turn gives β_2.

A.3.3. Converting short-horizon price impact estimates to a quarterly horizon

As a point of comparison, I linearly aggregate four existing short-horizon estimates of *total* price impact (temporary plus permanent) across a calendar quarter. Three of these estimates are from the academic literature, while one is an industry estimate. Reassuringly, my estimate of the permanent price impact component is less than these estimates of total price impact.[4]

Frazzini, Israel, and Moskowitz (2012) estimate that trading 1% of daily volume in a

[4] These calculations do not account for the fact that I measure trades relative to lagged volume, while the authors of these estimates measure trades relative to contemporaneous volume. In my sample, contemporaneous volume tends to increase relative to lagged volume for high volume consumed positions. Using contemporaneous volume, my estimates of permanent price impact would decline slightly.

U.S. equity generates 1.30 bps of market impact (their Table 5, column 8, the coefficient that describes price impact that is linear in the fraction of daily volume). Aggregating this figure suggests that trading 1% of volume for an entire quarter generates: 63 trading days *1.30 bps = 0.82% total price impact.

Collin-Dufresne and Fos (2015) find that on average, when 13D filers trade prior to their public filing date, they take up 26.1% of daily volume (their Table 1 row 10). On those same days, the excess return averages 34 bps (their Table 6 column 2). Thus trading 1% of daily volume generates an estimated market impact of (34 bps / 26.1% =) 1.30 bps, the same figure as in Frazzini, Israel, and Moskowitz (2012).

Brennan and Subrahmanyam (1996) find that the the average price impact generated by purchasing 1% of the shares outstanding of a stock in the middle quintile of size (market cap) and the middle quintile of illiquidity is 1.7% (their Table 1, panel B, estimates of C_n). In my sample, on average quarterly volume is roughly 50% of the market cap of a stock (my Table 1.1). Brennan and Subrahmanyam's estimate thus implies that trading 1% of volume for an entire quarter generates: 1.7% *50% = 0.85% of total price impact.

Investment Technology Group estimates a price impact of approximately 85 bps for consuming 5% of the volume in a $1.4 billion market cap stock over 30 days.[5] This estimate is based on the average execution price of an order (the weighted average of shares traded and the price of each transaction), so it represents a lower bound on the total price impact (final price minus initial price). Early trades will presumably be executed before prices have moved substantially. Nevertheless, aggregating ITG's estimate suggests that trading 1% of volume for an entire quarter generates at least: 85 bps /5% *3 months = 0.51% of total price impact. At the extreme, if one assumes that all price impact is permanent and that component trades are made in infinitesimally small amounts, then the price impact

[5] Hanson, Samuel G. "The FLV Capital Trading Desk (A)." Harvard Business School Teaching Note 215-053, January 2015.

after all trades have been executed will simply be twice this amount (1.02%). Thus ITG's comparable estimate of price impact most likely falls somewhere in the range of 0.51% to 1.02%, the midpoint of which is 0.77%.

A.4. Data appendix

A.4.1. Standardized earnings surprises (SUE)

In the data, to ensure that earnings reflect firm performance over the same period that hedge funds are trading, I include only companies with calendar quarter-end fiscal periods to match the 13F effective dates. For these companies, hedge fund trading over the course of quarter t can be mapped to an earnings release that reflects company performance over that same quarter.

The earnings return in quarter $t+1$ is measured as the return over the three trading-day window centered around the first Compustat earnings announcement date in quarter $t+1$ for stock s, using characteristic-adjusted daily returns. The return on "other days" in the quarter is the average daily characteristic-adjusted return outside of the earnings window, multiplied by three for comparability.

Standardized unexpected earnings (SUE), $SUE_{s,t}$, is measured as $\frac{earnings_{s,t+1} - median\ analyst\ forecast_{s,t}}{P_{s,t}}$, as in Baker, Litov, Wachter, and Wurgler (2010). To form the median forecast, I take the median across the last earnings forecast made by each analyst who published an earnings forecast during quarter t. I use only analyst forecasts made during quarter t to ensure that forecasts are made during the same time interval over which I measure hedge fund trades.

To faciliate interpretation, I standardize SUE in the cross section by quarter.

A.4.2. Constructing mutual fund flows

I identify funds subject to extreme fund-flows as in Coval and Stafford (2007, CS).

First, I link CRSP mutual fund returns and assets to the Thompson Reuters mutual fund holdings data, using the MFLINKS dataset provided by WRDS. As in CS, I remove funds with an IOC code (Thompson Reuters) of international, municipal bonds, bonds and preferred, or metals (1, 5, 6, or 8). I also eliminate funds with fewer than 5 holdings or with less than \$5 million in assets. I aggregate multiple share classes in CRSP (which are all linked to a single Thompson Reuters fund-quarter holdings entry), summing assets and forming returns as the asset-weighted average return of the underlying share classes. I then use the CRSP data to measure fund flows for fund f during quarter t: $FLOW_{f,t}^{crsp} = assets_{f,t}^{crsp} - assets_{f,t-1}^{crsp} * (1 + ret_{f,t}^{crsp}) - mergers_{f,t}^{crsp}$, where $ret_{f,t}^{crsp}$ is the return of fund f from the end of quarter $t-1$ until the end of quarter t, $assets_{f,t}^{crsp}$ is the total net assets of fund f at the end of quarter t, and $mergers_{f,t}^{crsp}$ represents the assets that fund f gained from mutual fund mergers during quarter t. I denote these variables as "CRSP" variables to explicitly denote that they are taken from CRSP, as opposed to returns and flows calculated using the holdings data (13Fs or mutual fund holdings). I then translate this into "relative" flows at the fund level: $flow_{f,t}^{crsp} = \frac{FLOW_{f,t}^{crsp}}{assets_{f,t-1}^{crsp}}$. I sort mutual funds into deciles at the end of each quarter t based on their $flow_{f,t}^{crsp}$.

Funds in the top decile of flows ($flow_{f,t}^{crsp}$) are "extreme inflow" funds, while funds in the bottom decile are "extreme outflow" funds.

A.4.3. Constructing funds' return gaps

Kacperczyk, Sialm, and Zheng (2008, KSZ) construct a measure of the differential between a fund's returns and the returns of its underlying holdings (assuming that trades are made costlessly at period ends), dubbed the fund's "return gap." KSZ find that funds with the highest return gaps (where the fund returns are much greater than the holding returns) generate the highest overall fund-level returns.

KSZ's Appendix A lists a comprehensive explanation of their sample selection. I follow the same process to identify mutual funds to include in the sample. KSZ filter by the Thompson Reuters (IOC) and CRSP (ICDI, Strategic Insight, Weisenberger, Policy) mutual fund objective codes. They also eliminate funds that hold less than 80% or above 105% in stocks, on average. KSZ eliminate funds with fewer than 10 holdings or with less than \$5 million in assets. They aggregate share classes in CRSP by forming the asset-weighted return of different shareclasses before matching with the Thompson Reuters holding data. I follow all of these procedures.

The monthly return gap is the differential between a fund's gross returns reported to CRSP (formed by taking the net return each month and adding back the expense ratio divided by 12) and the returns of the fund's most recently reported asset holdings during that month. With m indexing months, the net fund return is reported to CRSP, $ret^{crsp}_{f,m}$, as described above in Appendix A.4.2. $grossret^{CRSP}_{f,m} = ret^{crsp}_{f,m} + \frac{expense\,ratio_{f,m-1}}{12}$, where $expense\,ratio_{f,m-1}$ is the fund's most recently reported annual expense ratio as of the previous month end. $holdret_{f,m} = \frac{\sum_s ret_{s,m} * shares_{s,f,m-1} * P_{s,m-1}}{\sum_s shares_{s,f,m-1} * P_{s,m-1}}$, where $shares_{s,f,m-1}$ are the most recent shareholdings reported by manager f in stock s as of the end of the previous month $(m-1)$, $P_{s,m-1}$ is the price of stock s as of the most recent month end $(m-1)$, and $ret_{s,m}$ is the total return of stock s during month m. I include fund holdings that are up to six months old when calculating holding period returns.

At each calendar quarter end, I rank funds into quintiles based on $\sum_{k=1}^{12} grossret^{CRSP}_{f,m-k+1} - holdret_{f,m-k+1}$, where m indexes months. In my analysis in Section 1.5.2, because I analyze fund trades I only include funds that file consecutive quarter end holdings reports.

A.5. Additional results

A.5.1. Additional summary and portfolio tables

Table A.2. Summary statistics

This table displays additional summary statistics of volume consumed (aggregation method 1) 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. The value at each quarter t is calculated as the equal-weighted average across all stocks s in the corresponding decile portfolio at t. For manager statistics, before averaging across stocks a data point is calculated for each stock s as the equal-weighted average across all funds f who purchased stock s during quarter t. For stock-characteristic quintile averages, the value of a given characteristic for stock s is calculated as of the end of quarter $t-1$, to distinguish stock characteristics from the potential price impact of trades during quarter t. For quintiles, 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. Stocks below the 20th percentile of NYSE market cap have been removed.

Decile of volume consumed (t)	Avg mgr # positions	Avg mgr assets ($ MM)	Avg mgr age (quarters)	Avg stock size quintile	Avg stock book quintile	Avg stock momentum quintile
1	88.4	$478	21.4	3.61	2.71	3.07
2	86.2	$559	20.9	3.69	2.69	3.13
3	84.1	$602	20.8	3.66	2.69	3.14
4	83.0	$652	21.1	3.59	2.73	3.16
5	81.8	$708	21.2	3.50	2.75	3.19
6	81.0	$780	21.5	3.40	2.76	3.22
7	79.4	$821	21.4	3.28	2.78	3.23
8	78.1	$898	21.5	3.14	2.79	3.17
9	76.0	$1,085	21.9	2.97	2.82	3.16
10	72.5	$1,116	21.9	2.73	2.85	3.14

Table A.3. Contemporaneous performance

This table displays the contemporaneous market-adjusted and characteristic-adjusted monthly performance of calendar-time portfolios sorted into deciles based on volume consumed in quarter t by aggregation methods 1 (columns 1-2), 2 (column 3), and 3 (column 4). For comparison, I also display the monthly performance during quarter $t+1$ of portfolios of all stocks sorted by quarter t characteristic-adjusted performance (method †) and the monthly performance during quarter t of portfolios sorted by $valoftrade^{open}_{s,t} = shares\,traded_{s,t} * P_{s,t-1}$ (method ‡). 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.

Column:	(1)	(2)	(3)	(4)	(5)	(6)
Agg. method:	(1)	(1)	(2)	(3)	(†)	(‡)
Decile of volume consumed (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)
1	0.25%	0.14%	0.44%	0.29%	0.03%	0.72%
	[1.43]	[1.47]	[6.29]**	[2.38]**	[0.11]	[5.14]**
2	0.43%	0.27%	0.51%	0.03%	0.16%	0.95%
	[2.42]**	[2.93]**	[6.95]**	[0.34]	[1.32]	[8.71]**
3	0.49%	0.37%	0.51%	-0.15%	0.18%	0.99%
	[2.95]**	[4.25]**	[6.64]**	[-1.79]*	[2.37]**	[9.95]**
4	0.60%	0.46%	0.50%	-0.09%	0.16%	0.82%
	[3.39]**	[4.85]**	[6.57]**	[-0.99]	[2.78]**	[7.91]**
5	0.65%	0.49%	0.64%	-0.09%	0.04%	0.78%
	[3.76]**	[5.67]**	[8.46]**	[-0.93]	[0.68]	[8.32]**
6	0.93%	0.69%	0.73%	0.33%	0.09%	0.80%
	[4.95]**	[7.87]**	[9.11]**	[3.51]**	[1.59]	[9.31]**
7	0.95%	0.73%	0.90%	0.55%	-0.01%	0.65%
	[5.01]**	[7.63]**	[12.38]**	[6.69]**	[-0.20]	[6.97]**
8	0.93%	0.73%	1.03%	0.76%	-0.12%	0.52%
	[4.68]**	[7.15]**	[12.36]**	[7.80]**	[-1.90]*	[5.70]**
9	1.38%	1.18%	1.15%	0.92%	-0.07%	0.54%
	[6.44]**	[10.13]**	[13.74]**	[9.29]**	[-0.81]	[6.06]**
10	2.28%	2.07%	1.63%	1.75%	0.03%	0.21%
	[9.90]**	[13.69]**	[14.61]**	[14.63]**	[0.16]	[2.04]**
L/S (10-1)	2.04%	1.94%	1.18%	1.47%	0.00%	-0.52%
	[9.54]**	[10.44]**	[8.44]**	[13.05]**	[0.01]	[-3.00]**

Table A.4. Future trading

This table displays the volume consumed (% of quarterly volume) during quarter t+1 of calendar-time portfolios sorted into deciles based on volume consumed in t by aggregation methods 1 (columns 1-2), 2 (column 3), and 3 (column 4). For each portfolio, volume consumed in quarter $t + 1$ is calculated using the same aggregation method used to calculate volume consumed during quarter t. Calculations are based on 13F filings from 12/31/1989-9/30/2012. Positions are weighted equally. T-statistics of the long-short portfolios 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)
Agg. method:	(1)	(1)	(2)	(3)
Decile of volume consumed (t)	Volume consumed (t)	Volume consumed (t+1)	Volume consumed (t+1)	Volume consumed (t+1)
1	0.05%	1.01%	0.01%	-1.21%
2	0.18%	1.22%	0.03%	-0.33%
3	0.37%	1.37%	0.05%	-0.05%
4	0.63%	1.52%	0.08%	0.12%
5	0.99%	1.81%	0.12%	0.33%
6	1.52%	2.06%	0.17%	0.42%
7	2.33%	2.42%	0.27%	0.46%
8	3.72%	2.94%	0.43%	0.48%
9	6.62%	3.75%	0.82%	0.63%
10	17.63%	5.11%	2.90%	0.98%
L/S (10-1)	17.59%	4.09%	2.89%	2.19%
	[34.02]**	[22.17]**	[23.28]**	[17.45]**

Table A.5. SUE and earnings returns

This table displays additional results involving earnings announcement returns and standardized earnings surprises. The characteristic-adjusted earnings return measures the return of stock s during the three trading-day window centered around its first earnings announcement during quarter t+1. SUE is the standardized earnings surprise for stock s, normalized to have a cross-sectional standard deviation of one each quarter. $ME_{s,t}$, $V_{s,t-1}^{-1}$, $IOR_{s,t}$, 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. Panel A shows the coefficient on SUE by three groups of volume consumed (none or bottom quintile, the middle three quintiles, and the top quintile) in a Fama-MacBeth regression of the earnings return on SUE using observations with positive SUE. Panel B repeats the analysis of panel A using using observations with negative SUE.

Table A.5: (continued)

Panel A: Regression of characteristic-adjusted earnings returns (t+1) on positive SUE (t+1) by volume consumed groups

Volume consumed (t)	Coefficient on SUE (t+1)	[t-stat]
None or bottom quintile	2.52%	[7.93]**
Middle quintiles	2.08%	[8.73]**
Top quintile	1.50%	[3.62]**
Volume consumed (t)	**Constant**	**[t-stat]**
None or bottom quintile	omitted	
Middle quintiles	0.16%	[1.65]*
Top quintile	0.54%	[3.37]**
Controls	**Coefficient**	**[t-stat]**
$ME_{s,t}$	-0.31%	[-4.01]**
$V_{s,t-1}^{-1}$	-0.06%	[-0.83]
$IOR_{s,t}$	0.75%	[4.23]**
$BEME_{s,t}$	-0.24%	[-5.49]**
Test	**F-stat**	**p-value**
SUE coefficient: top-bottom?	6.35	0.014**
Fama-MacBeth		Y
Observations		80,362
R-squared		0.047

Table A.5: (continued)

Panel B: Regression of characteristic-adjusted earnings returns (t+1) on negative SUE (t+1) by volume consumed groups

Volume consumed (t)	Coefficient on SUE (t+1)	[t-stat]
None or bottom quintile	0.42%	[8.05]**
Middle quintiles	0.30%	[4.55]**
Top quintile	0.49%	[3.69]**
Volume consumed (t)	**Constant**	**[t-stat]**
None or bottom quintile	omitted	
Middle quintiles	0.21%	[2.49]**
Top quintile	0.89%	[5.53]**
Controls	**Coefficient**	**[t-stat]**
$ME_{s,t}$	0.35%	[3.87]**
$V_{s,t-1}^{-1}$	0.16%	[2.08]**
$IOR_{s,t}$	-1.10%	[-4.17]**
$BEME_{s,t}$	0.50%	[6.56]**
Test	**F-stat**	**p-value**
SUE coefficient: top-bottom?	0.24	0.627
Fama-MacBeth		Y
Observations		40,387
R-squared		0.047

Table A.6. Mutual fund trades, volume consumed, and performance

This table displays the volume consumed (% of quarterly volume) and monthly performance of mutual fund trades. Stocks are sorted into deciles based on volume consumed (aggregation method 1) during quarter t. Calculations are based on mutual fund holdings from 12/31/1989-9/30/2012 (except for active share results, which end at 12/31/2009). Positions are weighted equally. T-statistics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively. In calculating future performance, stocks below the 20th percentile of NYSE market cap have been removed. Volume consumed has been winsorized at the 1%/99% levels. Panel A displays volume consumed, contemporaneous performance, and future performance of mutual fund trades. Panel B displays future performance of the trades of subsets of mutual funds: funds in the top/bottom quintile of return gap or funds with above/below median active share.

Panel A: Mutual fund volume consumed – contemporaneous performance and future performance

Column:	(1)	(2)	(3)	(4)	(5)
Decile of volume consumed (t)	Volume consumed (t)	Mkt.-adj ret (t)	Char.-adj ret (t)	Mkt.-adj ret (t+1)	Char.-adj ret (t+1)
1	0.10%	0.22%	-0.01%	0.15%	0.03%
		[1.35]	[0.10]	[0.84]	[0.32]
2	0.31%	0.40%	0.22%	0.19%	0.10%
		[2.61]*	[2.31]**	[1.29]	[1.10]
3	0.56%	0.57%	0.36%	0.25%	0.05%
		[3.73]**	[4.39]**	[1.76]*	[0.55]
4	0.83%	0.67%	0.45%	0.11%	0.06%
		[4.53]**	[5.64]**	[0.83]	[0.72]
5	1.14%	0.81%	0.59%	0.14%	0.05%
		[5.19]**	[6.61]**	[1.13]	[0.66]
6	1.54%	0.85%	0.70%	0.20%	0.10%
		[5.65]**	[8.24]**	[1.56]	[1.36]
7	2.06%	0.91%	0.68%	0.21%	0.11%
		[5.68]**	[7.96]**	[1.66]*	[1.42]
8	2.82%	1.02%	0.79%	0.20%	0.07%
		[5.77]**	[7.96]**	[1.64]	[0.83]
9	4.20%	0.96%	0.74%	0.11%	-0.05%
		[5.11]**	[6.76]**	[0.93]	[-0.67]
10	7.48%	1.06%	0.79%	0.22%	0.17%
		[4.38]**	[5.28]**	[1.79]*	[1.85]*
L/S (10-1)	7.38%	1.05%	0.80%	0.07%	0.14%
		[5.45]**	[5.22]**	[0.46]	[1.08]

Table A.6: (continued)

Panel B: Mutual fund subsets – future performance

Column:	(1)	(2)	(3)	(4)
	Return gap, top quintile	Return gap, bottom quintile	Active share > median	Active share < median
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	0.20%	0.13%	0.03%	0.10%
	[1.75]*	[1.16]	[0.29]	[0.79]
2	0.03%	0.01%	0.10%	0.00%
	[0.28]	[0.13]	[0.93]	[-0.02]
3	-0.07%	0.12%	0.00%	0.05%
	[-0.71]	[1.20]	[-0.02]	[0.48]
4	0.01%	0.05%	0.21%	0.05%
	[0.09]	[0.61]	[2.04]**	[0.54]
5	0.06%	0.06%	-0.09%	0.03%
	[0.59]	[0.58]	[-1.08]	[0.36]
6	0.07%	0.03%	0.09%	-0.02%
	[0.73]	[0.41]	[0.91]	[-0.21]
7	0.25%	0.19%	-0.02%	0.14%
	[2.50]**	[2.16]**	[-0.20]	[1.89]*
8	0.26%	0.10%	0.39%	0.05%
	[2.49]**	[1.08]	[4.27]**	[0.56]
9	0.21%	0.23%	0.26%	0.09%
	[1.90]*	[2.32]**	[2.69]**	[1.09]
10	0.37%	0.07%	0.29%	-0.02%
	[3.15]**	[0.58]	[2.78]**	[-0.21]
L/S (10-1)	0.18%	-0.06%	0.26%	-0.12%
	[1.03]	[-0.42]	[1.62]	[-0.82]

A.5.2. Wurgler and Zhuravskaya (2002) and "best ideas"

Wurgler and Zhuravskaya (2002, WZ) also motivate weighted idiosyncratic risk as a trade-level limit, similar to Cohen, Polk, and Silli (2010). WZ model an arbitrageur that has exponential utility with constant absolute risk aversion (who is thus a mean-variance optimizer). When the arbitrageur is aware of a mispriced stock, she buys (or sells) that stock and attempts to hedge the position with a substitute portfolio. In this framework, idiosyncratic risk captures the risk of the trade after hedging. WZ show that variation in idiosyncratic risk helps explain cross-sectional variation in stock returns around index additions.

WZ use two empirical proxies for idiosyncratic risk. The first proxy is the variance of the simple market-adjusted return of a stock. The second proxy is the variance of a stock's return relative to the return of a characteristic-matched portfolio. The matching portfolio is constructed by finding three stocks in the same industry with similar market capitalizations and book-to-market ratios to the stock in question.

CAPM idiosyncratic variance – which I employ in Section 1.7.2 to identify funds' "best ideas" – closely corresponds to WZ's first proxy. It captures the risk remaining in a stock after the arbitrageur hedges that stock using the (beta-weighted) market portfolio. In unreported results, I also employ the variance of stocks' characteristic-adjusted returns to proxy for idiosyncratic risk when identifying funds' best ideas. This proxy is similar in spirit to WZ's second proxy. This measure of risk implicitly supposes the arbitrageur hedges her position in a stock with its characteristic-matched (DGTW) portfolio. Using characteristic-adjusted idiosyncratic risk produces similar results to my results using CAPM idiosyncratic risk. Best ideas remains uninformative.

Robustness to this variation is consistent with WZ. WZ find that the correlation between their two measures of idiosyncratic risk is 0.98. WZ find that idiosyncratic risk is difficult to hedge in general, as it is hard to find close substitutes for individual stocks.

A.5.3. Best ideas extended results

Table A.7. Volume consumed and best ideas

This table displays results comparing volume consumed and best ideas. It shows the characteristic-adjusted monthly performance during quarter $t+1$ of calendar-time portfolios sorted independently along measures of volume consumed and best ideas in quarter t. Positive volume consumed (aggregation method 2) positions are sorted into quintiles, with all positions with zero or negative values placed into a separate bin. Positions are independently sorted by their intra-manager best ideas ranking (relative to other stocks s in fund f's portfolio at quarter t). Calculations are based on 13F filings from 12/31/1989-9/30/2012. Positions are weighted equally. Stocks below the 20th percentile of NYSE market cap have been removed. T-statistics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively. The proportion of total positions within each bin is displayed in italics.

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

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

		21+	11-20	6-10	4-5	2-3	1
	Sale or hold	0.08%	0.08%	0.05%	0.06%	0.05%	0.06%
		[2.95]**	[1.81]*	[0.82]	[0.64]	[0.51]	[0.38]
		32.4%	*7.3%*	*3.8%*	*1.5%*	*1.5%*	*0.7%*
Volume consumed quintile (t)	1	0.03%	0.27%	0.02%	0.31%	0.11%	0.36%
		[0.78]	[2.96]**	[0.19]	[1.62]	[0.49]	[1.19]
		8.6%	*1.4%*	*0.7%*	*0.2%*	*0.2%*	*0.1%*
	2	0.04%	0.06%	-0.08%	0.02%	0.08%	0.04%
		[0.78]	[0.64]	-[0.55]	[0.08]	[0.38]	[0.16]
		7.5%	*1.7%*	*0.9%*	*0.3%*	*0.3%*	*0.1%*
	3	0.17%	0.07%	0.02%	-0.01%	-0.01%	0.28%
		[3.44]**	[0.80]	[0.17]	-[0.07]	-[0.03]	[1.01]
		6.5%	*2.1%*	*1.1%*	*0.4%*	*0.4%*	*0.2%*
	4	0.27%	0.19%	0.13%	0.09%	0.20%	-0.01%
		[5.40]**	[2.32]**	[1.24]	[0.59]	[1.27]	-[0.04]
		5.4%	*2.3%*	*1.3%*	*0.5%*	*0.5%*	*0.2%*
	5	0.30%	0.33%	0.57%	0.30%	0.30%	0.41%
		[4.91]**	[4.86]**	[6.27]**	[2.40]**	[2.20]**	[2.02]**
		3.9%	*2.4%*	*1.6%*	*0.7%*	*0.7%*	*0.4%*

A.6. Competition

A.6.1. Competition

The Kyle model assumes the insider is an information monopolist. The model can be extended to the case of multiple informed agents.

In my empirical results, I primarily focus on aggregating purchases at the stock-quarter level because price impact should aggregate. Disaggregated purchases at the stock-fund-quarter level also strongly predict future stock performance (Table 1.3). In this appendix, I examine variation in how multiple funds simultaneously trade a single stock.

Holden and Subrahmanyam (1992; HS) show that in a Kyle model with multiple identically informed agents and a large number of periods, informed traders aggressively compete, rapidly pushing prices towards fair value. Foster and Viswanathan (1996; FV) and Back, Cao, and Willard (2000) show that in contrast, prices gradually move towards fair value over time – as in the single agent case – if the informed agents' private signals are sufficiently heterogeneous.[6] These theoretical studies take the level of competition as exogenously fixed.

Applying these multi-agent versions of the Kyle model to my data poses several challenges. First, competition is not exogenously fixed. Competition varies based on how funds assign their limited attention. To a first approximation, competition may be randomly assigned. I model this extension below. In reality, competition is endogenous. Skilled funds may be adept at deciding what stocks to learn more about: more mispriced stocks may attract more competition. Second, in order to model competition, one must take a stance on the information structure underlying not only asset prices and what funds know about asset prices, but also what funds know about what other funds know about asset prices. Funds act based on their expectations of competitors' behavior. Third, the models assume that agents act independently. Some funds may coordinate their actions, as many hedge fund

[6] Koudijs (2014) also notes this distinction when applying the Kyle model to his data.

managers share common employment and educational backgrounds.[7] Fourth, the number of time periods has major implications for some competitive effects. My assumption that trade occurs once a quarter – the frequency of my data – is more stark in such an environment.

I explicitly elaborate on the first point. In Appendix A.6.2, I construct a one-period Kyle model that features a random level of competition. Each of two traders are randomly active or inactive. I assume the econometrician can only observe informed purchases, an aspect of my data. This model makes a key point: observing more insiders purchasing an asset increases estimated price impact but also increases the expected value of the asset. The correlation of the informed traders' signals determines which effect dominates.

If the econometrician observes a single informed trader purchasing a stock, the econometrician may expect that stock will perform particularly well in the future. The informed trader was able to build her position at a lower price because the second trader did not also purchase shares. However, observing two informed traders purchasing a stock increases the estimate of the asset's value. When forming the posterior distribution of the information, two observations receive more weight than one. Furthermore, the second informed trader may have been active but received a negative signal (and therefore gone unobserved).[8]

With perfectly correlated signals, observing a single informed trader leads to greater expected future returns. Signals are identical, so a second purchase would not increase the expected value of the asset. In contrast, with relatively uncorrelated signals, a second purchase increases the expected value of the asset by more than the incremental price impact. The same reasoning applies to observing a single trader purchasing $2x$ shares compared to observing two traders who each purchase $1x$ shares. I illustrate these points with a

---

[7] In perhaps the best known example, a number of proteges of Julian Robertson manage hedge funds. This group of funds, known as "Tiger Cubs," frequently trade in the same stocks. "There at least 30 'Tiger Cubs'...[and] 40-odd 'Tiger Seeds,' or funds that are backed by Robertson's money....it is believed that many of the managers still share ideas," from http://www.benzinga.com/trading-ideas/long-ideas/12/09/2876699/the-five-stocks-tiger-cubs-love#ixzz3n4IPTUFL, accessed 9/15/2015.

[8] If hedge funds endogenously allocate their attention, random assignment may understate this effect. Hedge funds may actively avoid competing with each other except in assets that are particularly mispriced.

parametrized example in Appendix A.6.4 and Figure A.3.

Empirically, when multiple funds simultaneously purchase a stock in my sample, the amount that each fund purchases varies substantially. When at least three funds purchase a stock simultaneously, the mean ratio of the standard deviation of volume consumed divided by average volume consumed is 1.25.[9] This pattern suggests information signals may be weakly correlated.[10]

In Table A.8, I present regressions of future returns on proxies for competition: the number of funds that purchase a stock (positively related to competition) and the average volume consumed in that stock (negatively related to competition). The number of funds that purchase a stock is positively related to the stock's future returns after controlling for its volume-consumed quintile. The average volume consumed is negatively related to future returns. However, the predictive effects of these proxies are insignificant when I limit the sample to the top quintile of volume consumed, where I have the strongest evidence that hedge funds trade based on information. These findings provide some evidence for the multi-agent Kyle model in which funds have relatively uncorrelated signals (as in FV). At the very least, more observable competition for a given total amount of trading does not appear to predict strongly diminished future returns (an implication of HS). The complications outlined above caution against interpreting these results too strongly.

A.6.2. Kyle model with a random number of informed traders (0, 1, or 2)

I construct a one-period Kyle model with an uncertain number of informed traders. Notation is the same as in Section 1.2, but without time subscripts. Each of two informed traders

[9] That is, I calculate $\frac{var(volconsumed_{s,f,t})_{s,t}^{1/2}}{\sum_{f=1}^{F} volconsumed_{s,f,t}/N_{s,t}}$ for stock s at quarter t. $N_{s,t}$ is the number of funds with positive volume consumed in stock s during quarter t, and the volatility calculation includes only positive observations of volume consumed in stock s during quarter t. I then average across stocks and quarters.

[10] This variation could also reflect non-information based motives for trade.

has a δ probability of being "active" in the asset, independent of whether the other trader is present. If both traders are active, they draw signals $i = \epsilon + \eta$ and $i' = \epsilon + \eta'$, with η and η' i.i.d. $N(0,\, \sigma_\eta^2)$. The market maker can not observe traders' presence, and therefore reacts to trades using a probability weighted average of the linear reaction function she would employ in each scenario.

"Future returns" are proxied by $\epsilon - p$. These are the returns realized by asset holders after trading takes place at price p.

The model solution proceeds as it does for the two-period model in Appendix A.1.

First, optimize from the perspective of an informed trader. She solves $max_x E[x(\epsilon - \lambda u - \lambda x - \delta\lambda\beta(\phi i')|\, i]$, where i is the agent's own information signal and i' is the information signal of the other agent (if that agent is active). The solution, after setting $\beta = \beta'$ (since the agents are identical), gives $\beta = \frac{1}{\lambda(2+\delta)}$.

Note that $E[p] = \lambda\beta\phi(i + i')$ if two traders are present, and $E[p] = \lambda\beta\phi i$ if one trader is present. Thus $\lambda\beta * (\#traders)$ represents the proportion of the informed traders' information that gets into prices in expectation. Suppose that at least a single informed trader is present. As $\delta \to 0$, this reduces to the classic Kyle model solution that $\lambda\beta = \frac{1}{2}$. If the single informed trader knows that the odds of her competing with another informed trader are approximately zero, then she will trade to get half of her information into price. As $\delta \to 1$, on the other hand, $\lambda\beta \to \frac{1}{3}$. Since both traders are active with certainty, that means that $2\beta\lambda = \frac{2}{3}$ of their information gets into price. Thus as more agents compete over the asset, they get more information into prices for a given true amount of information ϵ.

The market maker posts a single linear response coefficient λ (so that $p = \lambda(x + u)$). The market maker probabilistically averages her response function across the scenarios of no active informed traders, one active informed trader, and two active informed traders:
$\lambda = 2(1-\delta)\delta\frac{\beta\phi\sigma_\epsilon^2}{\beta^2\phi^2(\sigma_\epsilon^2+\sigma_\eta^2)+\sigma_u^2} + \delta^2\frac{\beta\phi\sigma_\epsilon^2}{4\beta^2\phi^2\sigma_\epsilon^2+2\beta^2\phi^2\sigma_\eta^2+\sigma_u^2}$.

A.6.3. Expected returns conditional on observing one vs. two traders when the econometrician only observes purchases

Assume the econometrician only observes informed purchases. In the model, if two informed traders purchase an asset, then the econometrician can infer the information of both traders. However, if the econometrician observes one informed purchase, she can not be sure if there was in fact only a single informed trader active or if instead a second informed trader was active but decided not to purchase (i.e., shorted) the asset.

In order to compare expected returns conditional upon observing informed purchases from one vs. two traders, three quantities are needed: (1) $E[\epsilon - p]$ if there is truly one trader active; (2) $E[\epsilon - p]$ if there are two traders active but one trader decides to short the asset; and (3) $E[\epsilon - p]$ if there are two traders active and both purchase the asset.

For the first quantity, the calculation is simple: $E[\epsilon - p|\, i] = E[\epsilon - \lambda\beta\phi i|\, i] = \phi i(\frac{1+\delta}{2+\delta})$, if only one trader is active and we observe i.

For the second quantity, we need the expectation of ϵ conditional on the second agent drawing a negative signal. This calculation utilizes the truncated normal distribution, so there is no analytical solution. Instead, solve for ϵ by maximizing its likelihood: $g(\frac{i-\epsilon}{\sigma_\eta})G(\frac{-\epsilon}{\sigma_\eta})g(\frac{\epsilon}{\sigma_\epsilon})$, with g and G the standard normal PDF and CDF, respectively (the first term represents the probability that the first informed trader draws a signal i, given ϵ; the second term reflects the probability that the second signal i' is negative, given ϵ; and the third term represents the prior probability that ϵ takes the given value). Then calculate the expected signal for the second trader, conditional on it being less than zero, using the moments of a truncated normal distribution (truncated at zero, with mean ϕi and variance $(1-\phi)\sigma_\epsilon^2 + \sigma_\eta^2$). Given the expected signal for the second trader, the expected value of ϵ, and ϕi for the observed trader, calculate $E[\epsilon - p] = E[\epsilon] - \lambda\beta(i + E[i'])$, with all the expectations conditional on observing i and knowing that the second unobserved trader receives a negative signal i'.

The third quantity is calculated similarly. We need the expectation of ϵ conditional on a hypothetical positive draw for the second trader. Solve for the expected value of ϵ by maximizing the likelihood $g(\frac{i-\epsilon}{\sigma_\eta})(1-G(\frac{-\epsilon}{\sigma_\eta}))g(\frac{\epsilon}{\sigma_\epsilon})$. Then proceed as above.

Finally, calculate the expected return conditional on observing one trader purchasing the asset as the probability weighted average of (1) there being only a single active trader and (2) the possiblity that a second trader was active but decided not to purchase the asset:

$$E[\epsilon - p_1 | \, observe\, one\, trader] = \frac{Prob(truly\ one\ trader)}{Prob(truly\ one\ trader) + Prob(unobserved\ second\ trader)} * E[\epsilon - p_1 | \, truly\, one\, trader] + \frac{Prob(unobserved\ second\ trader)}{Prob(truly\ one\ trader) + Prob(unobserved\ second\ trader)} * E[\epsilon - p_1 | \, unobserved\, second\, trader].$$

Compare that quantity to $E[\epsilon - p_1 | \, observe\, two\, traders]$.

An alternative manner of conceptualizing this dynamic is to consider expected returns conditional on observing a single purchase of $2x$ to expected returns conditonal on observing two smaller purchases that sum to $2x$.

The latter expectation is trivial. Assuming that the two traders observe signals of i_{two} and i'_{two}, where $i_{two} + i'_{two} = i_{one}$, with i_{one} the signal of the single large trader and $i_{two}, i'_{two}, i_{one} > 0$, then $E[\epsilon - p] = E[\epsilon] - \lambda\beta\phi(i_{two} + i'_{two})$. $E[\epsilon] = \frac{\frac{2}{\sigma_\eta^2}(i_{two} + i'_{two})/2}{\frac{1}{\sigma_\epsilon^2} + \frac{2}{\sigma_\eta^2}}$, based on forming a normal posterior from a prior (the distribution of ϵ) and data (observations of i and i').

The former expectation is calculated using the same method as above: probabilistically average the expectation if the trader is active on her own and the expectation if a second trader was active but decided not to purchase the asset.

A.6.4. Parametrized example

To get a quantitative sense of these dynamics, I construct a parametrized example of the model. I assume the econometrician only observes informed purchases.

Figure A.3 illustrates expected returns, $E[\epsilon - p]$, as one varies the noise of the informed traders' signals, σ_η^2. I assume that $\sigma_\epsilon^2 = 1$ and $\delta = 0.5$ (note that I do not need to make an assumption on σ_u^2, since I only need to know $\lambda\beta$, not λ on its own). These results are based on 50,000 simulations of the model for each value of σ_η^2. In each simulation, I randomly draw a positive value of $i = \epsilon + \eta$. I then calculate expected returns conditional on seeing a single purchase based on that signal. I also calculate expected returns if one were to observe a (random) second informed purchase. Finally, I calculate expected returns if instead of seeing a single informed purchase, the econometrician observes two informed purchases that are each half the size of the (larger) single purchase.

Expected returns are higher conditional upon observing a single informed purchase, compared to what one would expect if one observed a second informed purchase, for $\sigma_\eta^2 \leq 1$. At higher values of σ_η^2, the increase in the expected value of ϵ from observing a second purchase outweighs the increase in price impact (expected value of p). In that part of the parameter space, expected returns are higher if the econometrician observes two informed purchases.

I also compare one informed purchase to two informed purchases that are each half the size of the single purchase.[11] In this scenario, the point of preference shifts to a higher value of σ_η^2. Expected returns are higher for observing a single informed purchase if $\sigma_\eta^2 \leq 4$. At higher values of σ_η^2, returns are higher conditional on observing two smaller informed purchases.

With random assignment of informed traders and an inability to observe shorts, expected returns may be higher after observing more purchases or after observing fewer purchases. The noise in informed traders' signals determines the relative ranking.

[11] Mathematically, this solution applies to any two purchases that add up to the magnitude of the single larger purchase.

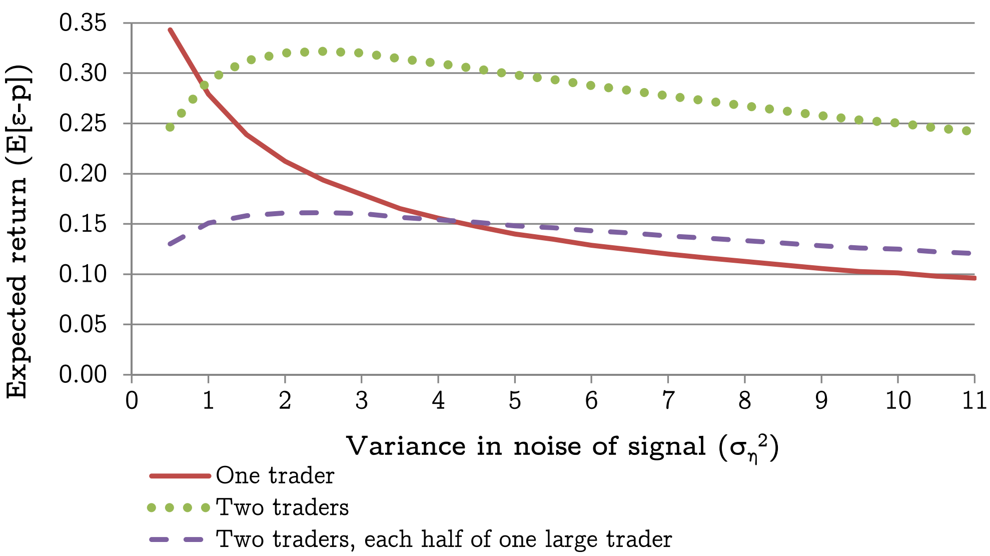

Figure A.3. Competition - parametrized example

This figure displays expected returns, $E[\epsilon - p]$, conditional on different observed patterns of trading in a one-period Kyle model with two randomly assigned informed traders, as in Appendix A.6.2. Expected returns are displayed as a function of σ_η^2. $\sigma_\epsilon^2 = 1$ (the variance of information) and $\delta = 0.5$ (the probability a given informed trader is "active" in a stock).

A.6.5. Competition – Results

Table A.8. Competition and future returns

This table displays information involving competition and future monthly characteristic-adjusted returns. $\#funds_{s,t}$ is the number of hedge funds that purchased a stock s in quarter t. The sample is limited to stocks with $\#funds_{s,t} > 0$. Average volume consumed is the average volume consumed in that stock: $\frac{volconsumed_{s,t}}{\#funds_{s,t}}$. VCQ is the quintile of volume consumed (aggregation method 1; 1-5 for stocks with hedge fund purchases, and 0 for stocks with no hedge fund purchases) for stock s. $ME_{s,t}$, $V_{s,t}^{-1}$, $IOR_{s,t}$, 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. 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)	Char.-adj ret (t+1)	Char.-adj ret (t+1)
VCQ (t)	0.0010	0.0015		
	[3.77]**	[5.21]**		
$\#funds_{s,t}$	0.0004		0.0002	
	[2.53]**		[0.68]	
Average volume consumed (t)		-0.0341		-0.0187
		[-2.35]**		[-1.07]
$ME_{s,t}$	-0.0009	-0.0007	-0.0031	-0.0031
	[-0.75]	[-0.61]	[-2.58]**	[-2.60]**
$V_{s,t-1}^{-1}$	0.0004	0.0003	0.0024	0.0022
	[0.38]	[0.33]	[2.60]**	[2.37]**
$IOR_{s,t}$	0.0022	0.0015	0.0053	0.0046
	[0.91]	[0.59]	[1.78]*	[1.63]
$BEME_{s,t}$	0.0012	0.0012	0.0006	0.0007
	[2.57]**	[2.60]**	[0.71]	[0.86]
Fama-MacBeth	Y	Y	Y	Y
Only top quintile of volume consumed?	-	-	Y	Y
Observations	148,996	148,996	30,278	30,278
R-squared	0.024	0.024	0.033	0.033

B. Appendix to Chapter 2

B.1. Additional results

Table B.1. Double down portfolio, without removing bottom decile of asset changes

This table displays the monthly performance of the double down portfolio, formed as described in the text, but without removing positions in which managers double down over a portfolio formation period during which the manager's proportional change in 13F assets falls in the bottom decile of the proportional change in assets of all managers in the sample. The baseline parameter values are used here. Portfolio performance is calculated from 12/31/1989-12/31/2013 for alpha calculations, and from 12/31/1989-6/30/2012 for DGTW calculations. Positions are weighted equally across managers but value-weighted within a given manager's portfolio, as described in the text. T-statstics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively.

Trailing ret interval	Raw return	DGTW adjusted	4 factor alpha	mkt	size	book	mom	CAPM alpha	mkt	Avg # positions in portfolio
3 mo	1.22%	0.41%	0.81%	1.09	-0.03	0.23	-0.31	0.58%	1.23	13.2
		[1.41]	[2.52]**	[14.5]	-[0.3]	[1.8]	-[2.7]	[1.89]*	[17.8]	
6 mo	1.52%	0.65%	0.64%	1.12	0.16	0.36	-0.26	0.54%	1.23	24.5
		[2.93]**	[2.84]**	[17.9]	[1.5]	[3.9]	-[4.4]	[2.25]**	[19.3]	
9 mo	1.39%	0.62%	0.56%	1.12	0.20	0.36	-0.26	0.47%	1.22	26.0
		[2.52]**	[2.12]**	[16.1]	[2.1]	[3.5]	-[3.0]	[1.80]*	[16.9]	
12 mo	1.21%	0.35%	0.33%	1.11	0.29	0.38	-0.18	0.33%	1.18	21.7
		[1.54]	[1.43]	[17.1]	[2.7]	[5.2]	-[2.8]	[1.39]	[17.0]	

Table B.2. Double down portfolio, different position cutoffs

This table displays the monthly performance of the double down portfolio, formed as described in the text, but defining sizable as the maximum of (1) X% and (2) the manager's average position size across all 13Fs filed by the manager to date. The value of X used for each portfolio is denoted in the table. Otherwise, the baseline parameter values are used. Portfolio performance is calculated from 12/31/1989-12/31/2013 for alpha calculations, and from 12/31/1989-6/30/2012 for DGTW calculations. Positions are weighted equally across managers but value-weighted within a given manager's portfolio, as described in the text. T-statstics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively.

Position	Dgtw-adjusted			4-factor alpha			CAPM alpha		
cutoff	3 mo	6 mo	9 mo	3 mo	6 mo	9 mo	3 mo	6 mo	9 mo
1.5%	0.43%	0.41%	0.42%	0.54%	0.41%	0.42%	0.35%	0.39%	0.45%
	[1.59]	[2.70]**	[2.48]**	[1.74]*	[2.57]**	[2.24]**	[1.10]	[2.33]**	[2.24]**
2.0%	0.53%	0.66%	0.55%	0.83%	0.57%	0.46%	0.62%	0.54%	0.50%
	[1.76]*	[3.69]**	[2.47]**	[2.59]**	[2.99]**	[1.91]*	[1.93]*	[2.79]**	[2.03]**
3.0%	0.02%	0.65%	0.72%	0.49%	0.81%	0.61%	0.25%	0.65%	0.63%
	[0.04]	[2.42]**	[2.61]**	[1.07]	[2.95]**	[2.22]**	[0.55]	[2.38]**	[2.26]**
3.5%	0.28%	0.80%	1.16%	0.86%	1.06%	1.05%	0.69%	0.88%	1.16%
	[0.54]	[2.16]**	[3.13]**	[1.75]*	[2.88]**	[3.00]**	[1.43]	[2.38]**	[3.23]**

Table B.3. Double down portfolio, different average factors

This table displays the monthly performance of the double down portfolio, formed as described in the text, but defining sizable as the maximum of (1) 2.5% and (2) X * the manager's average position size across all 13Fs filed by the manager to date. The value of X used for each portfolio is denoted in the table. Otherwise, the baseline parameter values are used. Portfolio performance is calculated from 12/31/1989-12/31/2013 for alpha calculations, and from 12/31/1989-6/30/2012 for DGTW calculations. Positions are weighted equally across managers but value-weighted within a given manager's portfolio, as described in the text. T-statstics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively.

Avg pos	Dgtw-adjusted			4-factor alpha			CAPM alpha		
factor	3 mo	6 mo	9 mo	3 mo	6 mo	9 mo	3 mo	6 mo	9 mo
0.50	0.44%	0.73%	0.61%	0.74%	0.62%	0.46%	0.47%	0.52%	0.47%
	[1.37]	[3.26]**	[2.70]**	[2.07]**	[2.96]**	[1.82]*	[1.35]	[2.34]**	[1.90]*
0.75	0.43%	0.75%	0.59%	0.70%	0.59%	0.42%	0.45%	0.52%	0.45%
	[1.29]	[3.38]**	[2.59]**	[1.92]*	[2.87]**	[1.69]*	[1.26]	[2.34]**	[1.80]*
1.25	0.37%	0.78%	0.69%	0.70%	0.76%	0.58%	0.51%	0.66%	0.55%
	[1.00]	[3.14]**	[2.69]**	[1.77]*	[3.07]**	[2.12]**	[1.33]	[2.60]**	[2.01]**
1.50	0.52%	0.87%	0.97%	0.76%	0.94%	0.83%	0.60%	0.88%	0.84%
	[1.30]	[2.97]**	[3.20]**	[1.86]*	[3.16]**	[2.62]**	[1.49]	[2.92]**	[2.65]**
2.00	0.25%	1.19%	1.06%	0.72%	1.29%	1.02%	0.52%	1.22%	1.07%
	[0.43]	[3.32]**	[2.97]**	[1.30]	[3.64]**	[2.77]**	[0.98]	[3.47]**	[2.96]**

Table B.4. Double down portfolio, different stock underperformance cutoff

This table displays the monthly performance of the double down portfolio, formed as described in the text, but requiring that over the relevant portfolio formation period, a stock's return must fall short of the CRSP value weighted market index by at least X%. The value of X used for each portfolio is denoted in the table. Otherwise, the baseline parameter values are used. Portfolio performance is calculated from 12/31/1989-12/31/2013 for alpha calculations, and from 12/31/1989-6/30/2012 for DGTW calculations. Positions are weighted equally across managers but value-weighted within a given manager's portfolio, as described in the text. T-statstics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively

Fall relative to mkt	Dgtw-adjusted			4-factor alpha			CAPM alpha		
	3 mo	6 mo	9 mo	3 mo	6 mo	9 mo	3 mo	6 mo	9 mo
0%	0.32%	0.45%	0.45%	0.55%	0.38%	0.32%	0.40%	0.33%	0.33%
	[1.65]*	[2.55]**	[2.21]**	[2.54]**	[2.06]**	[1.64]	[1.97]**	[1.75]*	[1.58]
5%	0.30%	0.57%	0.56%	0.53%	0.46%	0.42%	0.36%	0.41%	0.36%
	[1.09]	[2.97]**	[2.62]**	[1.86]*	[2.28]**	[2.02]**	[1.26]	[2.00]**	[1.68]*
15%	0.96%	1.17%	0.72%	1.40%	1.01%	0.48%	1.22%	0.99%	0.50%
	[2.00]**	[3.64]**	[2.47]**	[2.76]**	[3.30]**	[1.72]*	[2.45]**	[3.13]**	[1.76]*

Table B.5. Double down portfolio, different position cutoffs

This table displays the monthly performance of the double down portfolio, formed as described in the text, but requiring that over the relevant portfolio formation period, the manager must have increased the position portfolio weight to X * its weight at the begining of the formation period. The value of X used for each portfolio is denoted in the table. Otherwise, the baseline parameter values are used. Portfolio performance is calculated from 12/31/1989-12/31/2013 for alpha calculations, and from 12/31/1989-6/30/2012 for DGTW calculations. Positions are weighted equally across managers but value-weighted within a given manager's portfolio, as described in the text. T-statstics are displayed in brackets. ** and * denote significance at the 5% and 10% levels, respectively.

Increase in port wtd	Dgtw-adjusted			4-factor alpha			CAPM alpha		
	3 mo	6 mo	9 mo	3 mo	6 mo	9 mo	3 mo	6 mo	9 mo
1.50	0.10%	0.32%	0.39%	0.40%	0.29%	0.38%	0.23%	0.29%	0.44%
	[0.56]	[2.29]**	[2.59]**	[1.93]*	[1.87]*	[2.41]**	[1.09]	[1.73]*	[2.47]**
1.75	0.41%	0.58%	0.46%	0.80%	0.47%	0.48%	0.62%	0.41%	0.49%
	[1.53]	[2.97]**	[2.62]**	[2.85]**	[2.51]**	[2.45]**	[2.27]**	[2.07]**	[2.41]**
2.25	0.29%	1.23%	0.80%	0.86%	1.20%	0.72%	0.43%	0.96%	0.71%
	[0.67]	[3.45]**	[2.83]**	[1.79]*	[3.55]**	[2.54]**	[0.94]	[2.75]**	[2.42]**

C. Appendix to Chapter 3

C.1. Portfolio construction

C.1.1. Expert portfolio

I construct the return of the expert portfolio, $R_{expert,t+1}$, as follows. Let f=1 to F index all funds and s=1 to S index all stocks. Fund manager f's portfolio weight in stock s as a percent of her long equity holdings at the end of quarter t is $\bar{w}_{s,f,t}$ (so $\sum_{i=1}^{I_k} \bar{w}_{s,f,t} = 1$ if fund f files a 13F in quarter t, or $\bar{w}_{s,f,t} = 0$ if fund f does not file a 13F in quarter t). $1_{\{s,f,t\}}^{expert}$ is an indicator variable set equal to 1 if stock s is an expert position for manager f at time t, and zero otherwise, and $w_{s,f,t} = \bar{w}_{s,f,t} / \sum_{f=1}^{F} \sum_{s=1}^{S} \bar{w}_{s,f,t} 1_{\{s,f,t\}}$ so that portfolio weights sum to 1.

$$R_{expert,t+1} = \sum_{f=1}^{F} \sum_{s=1}^{S} w_{s,f,t} 1_{\{s,f,t\}}^{expert} r_{s,t+1}$$

C.1.2. Expert portfolios separated by trailing stock-level performance

I construct $R_{expert\,after\,poor\,trailing\,performance,t+1}$ as follows. I add an indicator variable $1_{\{s,f,t,t_{first\,expert}\}}^{poor\,trailing\,performance}$ that is set equal to 1 if stock s is an expert position for fund f at time t and the trailing return of stock i over the desired time interval (6, 9, or 12 months) was below the median stock return (i.e., the median stock return of all listed U.S. common stocks) in my sample over the time interval ending at $t_{first\,expert}$, where $t_{first\,expert}$ is the first quarter end for which stock s is an expert position for fund f such that it is also an expert position for all quarter ends $t_{interim}$ such that $t_{first\,expert} \leq t_{interim} \leq t$,[1] and zero otherwise. I set $\hat{w} = \bar{w}_{k,i,t} / \sum_{k=1}^{K} \sum_{i=1}^{I_k} \bar{w}_{k,i,t} 1^{poor\,trailing\,performance}_{\{k,i,t,t_{first\,expert}\}}$ so that the portfolio weights sum to 1.

$$R_{expert\,after\,poor\,trailing\,performance,t+1} = \sum_{f=1}^{F} \sum_{s=1}^{S} \hat{w}_{s,f,t} 1^{poor\,trailing\,performance}_{\{s,f,t,t_{first\,expert}\}} r_{s,t+1}$$

I similarly construct $R_{expert\,after\,good\,trailing\,performance,t+1}$ as the returns of the portfolio of all expert positions that were not initiated after poor stock level performance.

C.1.3. Control portfolios

I construct the control portfolio of all positions that hedge fund managers purchase after poor trailing performance as follows. I label this portfolio's returns $R_{all\,positions\,after\,poor\,trailing\,performance}$. The indicator variable $1^{poor\,trailing\,performance,\,all\,positions}_{\{s,f,t,t_{first\,initiated}\}}$ is set to 1 if stock s is initiated in fund f's portfolio following poor trailing performance, and zero otherwise. I set $\hat{\hat{w}} = \bar{w}_{s,f,t} / \sum_{f=1}^{F} \sum_{s=1}^{S} \bar{w}_{s,f,t} 1^{poor\,trailing\,performance,\,all\,positions}_{\{s,f,t,t_{first\,initiated}\}}$ so that portfolio weights sum to 1.

$$R_{all\,positions\,after\,poor\,trailing\,performance,t+1} = \sum_{f=1}^{F} \sum_{s=1}^{S} \hat{\hat{w}}_{s,f,t} 1^{poor\,trailing\,performance,\,all\,positions}_{\{s,f,t,t_{first\,initiated}\}} r_{s,t+1}$$

I construct the characteristic-matched potential reversal control portfolio as follows. $mktcap_{s,t}$ equals the market cap of stock s at time t, $l_{s,t}$ equals the size quintile of stock

[1] In other words, $t_{first\,expert}$ represents the quarter end at which fund f initiated the current expert position in stock s. If fund f exits that expert positions at t_{exit} and subsequently reenters it at a time $t_2 > t_{exit}$, then $t_{first\,expert}$ will be set to the reinitation date for subsequent time periods $t \geq t_2$ during which stock s is an expert position for manager f.

s at time t, and $m_{s,t}$ equals the book-to-market quintile of stock s at time t. $s^{l_{s,t},m_{s,t}}$ indexes the set of $I^{l_{s,t},m_{s,t}}$ stocks in the $l_{s,t}$ size quintile and $m_{s,t}$ book-to-market quintile. $1_{\{s,t\}}^{poor\,trailing\,performance,\,all\,positions}$ is now an indicator variable that is set to 1 if the trailing performance of stock i ending at quarter t fell short of the median stock over the desired trailing time interval (6, 9, or 12 months), H equals the holding period for reversal positions (so $H = 3$ if positions are held for 3 quarters, or 9 months, after a dip). I set

$$w_{i^{l_{s,t},m_{s,t}},t}^{reversal\,portfolio} = \frac{mktcap_{s^{l_{i,t},m_{i,t}},t}}{\sum_{s'^{l_{s,t},m_{s,t}}=1}^{S'^{l_{s,t},m_{s,t}}}(\sum_{t'=t-H+1}^{t} 1_{\{s'^{l_{s,t},m_{s,t}},t'\}}^{poor\,trailing\,performance,\,all\,positions})mktcap_{s'^{l_{s,t},m_{s,t}},t}}$$

so that the characteristic matched portfolio weights sum to 1. The return of each of these 25 subportfolios indexed by $l_{i,t}, m_{i,t}$ is :

$$R_{t+1}^{l_{i,t},m_{i,t}} = \sum_{i^{l_{i,t},m_{i,t}}=1}^{I^{l_{i,t},m_{i,t}}} r_{i^{l_{i,t},m_{i,t}},t+1} w_{i^{l_{i,t},m_{i,t}},t}^{reversal\,portfolio}$$

Once these 25 size and value reversal portfolios have been formed, I construct the control portfolio by weighting the 25 subportfolios each quarter to match the weights that the portfolios constructed in panel A of Table 3.4 place on each of the corresponding 25 size and value buckets.

$$R_{characterisic\,matched,t+1} = \sum_{k=1}^{K}\sum_{i=1}^{I_k} w_{k,i,t} 1_{\{k,i,t,t_{first\,expert}\}}^{poor\,trailing\,performance} R_{t+1}^{l_{i,t},m_{i,t}}$$

C.2. Available expert capital

I define available expert capital as follows. For stock s at time t, where $1_{\{s,f,t\}}^{expert-eligible}$ is an indicator variable set equal to 1 if stock s is expert-eligible for fund f at time t and 0 otherwise, F is the set of managers in the sample, $mktcap_{s,t}$ is the market cap of stock s at time t, and $AUM_{f,t}$ is the total value of all stock positions in fund f's 13F filing at time t, $available\,expert\,capital_{s,t}$ is defined as:

$$available\, expert\, capital_{s,t} = \sum_{f=1}^{F} 1_{\{s,f,t\}}^{expert\, eligible} AUM_{f,t}/mktcap_{s,t}$$