Hidden Risk*

Daniel Barth† Phillip J. Monin‡ Emil Siriwardane§

Adi Sunderam¶

Abstract

This paper investigates the agency problems that arise when investors cannot observe the risks taken by their asset managers. Using a unique regulatory dataset covering virtually all U.S. hedge funds, we find that the scope for agency problems is large, as managers possess valuable information about fund risk beyond what investors can infer from past returns. Consistent with strategic communication, investor flows respond to this information only when it serves the managers' best interests. We estimate that 10-20% of hedge fund flows over the past decade may have been allocated differently had investors better understood fund risk taking. Such agency frictions are likely more severe in other parts of asset management, where risks are harder to infer from past returns.

*The views stated herein are those of the authors and are not necessarily the views of the Federal Reserve Board or the Federal Reserve System. We are grateful to seminar participants at the Federal Reserve Board, NYU Stern, and Stanford GSB for helpful comments.

†Barth: Board of Governors of the Federal Reserve System. E-mail: daniel.j.barth@frb.gov

‡Monin: Board of Governors of the Federal Reserve System. E-mail: phillip.monin@frb.gov

§Siriwardane: Harvard Business School and NBER. E-mail: esiriwardane@hbs.edu

¶Sunderam: Harvard Business School and NBER. E-mail: asunderam@hbs.edu

1 Introduction

Financial intermediation in the United States has changed dramatically over the last forty years. While traditional banks and mutual funds once dominated the investment landscape, intermediation chains are now longer and take place through less-regulated vehicles like hedge funds and private equity funds (Greenwood et al., 2024). At a more fundamental level, intermediation has shifted towards more arms-length transactions, with investors increasingly entrusting their capital to external fund managers operating under a variety of complex contractual arrangements (Rajan, 2005).

The benefits of this evolution are clear. With deeper and more efficient capital markets, public and private firms can raise financing from a broader set of investors at lower cost. By expanding households' access to financing and increasing opportunities for diversification, the proliferation of asset management also facilitates risk-sharing at a global scale. Moreover, the move away from relationship-based financing can curb cronyism and favoritism, allowing investment to flow more readily into productive enterprises.

Nevertheless, the increased complexity of financial intermediation also carries costs. Multiple layers of financial intermediation—coupled with less transparent investment vehicles—can hinder investors' ability to monitor the risks taken by their asset managers (Rajan, 2005, Acharya et al., 2016). This in turn complicates the problem of investors trying to assess the skill of fund managers. Investors cannot be sure if high realized returns result from risk taking or skill, particularly when managers invest in assets with low liquidity or option-like payoffs (Lo, 2001, Jurek and Stafford, 2015, Stafford, 2021). Thus, from the perspective of a fund manager, there are often strong incentives to take risks that investors do not recognize.

Although the potential for such agency problems in the modern financial system is widely acknowledged, gauging their quantitative significance is much more difficult. The primary challenge lies in measurement: in most datasets, it is not possible to determine the scale of the informational asymmetry between asset managers and their investors.

Moreover, holding fixed the size of the asymmetry, it also necessary to measure the degree to which contracting, reputation, and other solutions mitigate the problem. In some agency models, the private information of agents is ultimately fully revealed in equilibrium (e.g., Stein, 1989, Milgrom, 1981). In others, some information is concealed, either because there is a pooling equilibrium or because principals are not sophisticated.

In this paper, we make progress on measuring agency problems in asset management using regulatory data from the Securities and Exchange Commission (SEC) Form PF. These data encompass the vast majority of the U.S. hedge fund industry, which managed $6 trillion of net assets and $11 trillion in gross assets as of the end of 2019 (Barth et al., 2023). Form PF is uniquely suited for studying agency frictions because it requires hedge fund managers to report their own perceived exposures to various risk factors (e.g., public equity markets) directly to the SEC. Crucially, however, the SEC does not share this information with investors. This feature allows us to compare managers' private assessments of risk with estimates that are readily available to investors, namely those implied by past fund returns.

Throughout our analysis, risk is defined using CAPM betas with respect to public equity markets. On Form PF, funds must report how their portfolio values will be affected by changes in equity prices of -20%, -5%, 5%, and 20%, which we use to infer managers' perceived CAPM betas.[1] The difference between manager-perceived betas (β^R) and those implied by historical returns ($\hat{\beta}$) is what we label as the beta gap. Although the average beta gap is zero, some managers consistently report more equity exposure to the SEC than historical returns suggest, while others consistently report less. This variation is substantial, as the standard deviation of the beta gap is over 0.3. For an equity premium of 10%—roughly the realized annual premium during our sample—a beta gap of 0.3 implies that managers perceive the annual CAPM alpha of their fund to be 300 bps lower than what is implied by historical returns.

For a non-zero beta gap to imply an agency problem, managers' beliefs must also be informative about risk. Using forecasting regressions, we find that manager-perceived betas

[1]Section 3.1.1 contains a more complete description of how we extract manager beliefs from Form PF.

consistently predict future equity market exposure in both the cross-section and the time series, even after controlling for CAPM betas estimated from historical returns. The scope for agency problems is therefore potentially large in our setting, since managers possess useful information about risk that investors could not easily infer from past returns.

Next, we develop a test for whether investors incorporate this information into their asset allocation decisions. In the spirit of Barber et al. (2016) and Berk and van Binsbergen (2016), we take a revealed preference approach that links net flows into a fund in quarter $t+1$ ($flow_{t+1}$) to the CAPM alpha estimated from its recent returns ($\hat{\alpha}_t$). Our innovation is to augment standard flow-performance regressions with an interaction term—the product of the beta gap ($\beta_t^R - \hat{\beta}_t$) and the recent market return (μ_t^M)—that captures differences between the manager's and investors' perceptions of the manager's skill:

$$flow_{t+1} = a + \kappa_1 \hat{\alpha}_t + \kappa_2 (\beta_t^R - \hat{\beta}_t)\mu_t^M + \varepsilon_{t+1}.$$

To see the intuition of this regression test, consider a fund with a positive beta gap during a period of strong market returns, meaning $(\beta_t^R - \hat{\beta}_t)\mu_t^M > 0$. Because the manager perceives more equity exposure than implied by historical returns ($\beta_t^R > \hat{\beta}_t$), the fund's realized CAPM alpha $\hat{\alpha}_t$ will be overstated from the perspective of the manager. If investors incorporate the manager's private risk assessment into their allocation decisions, flows should be lower than predicted by the fund's historical alpha alone. In the performance-flow regression above, this means we should observe $\kappa_2 < 0$.

Across all funds, we find only weak evidence in support of this prediction, but the average effect masks significant heterogeneity in investor behavior. In particular, investors seem to respond strongly to manager risk perceptions when the beta gap is negative, yet show no such reaction in funds with positive beta gaps. In other words, when a fund's historical alpha exceeds the manager's perceived alpha, manager beliefs do not appear to influence investor flows. Similarly, investors do not respond to manager assessments of risk in funds with option-like equity exposure. Since fees are similar across these fund groups,

differences in flows are unlikely to be driven by investment costs.

While we cannot definitively pin down the microfoundations of these results, they are consistent with strategic communication by fund managers. Managers with positive beta gaps perceive lower risk-adjusted performance than implied by their returns. Thus, communicating their perceptions to investors could make these managers worse off by reducing flows and hence assets under management. Conversely, managers with negative beta gaps perceive higher risk-adjusted performance than implied by their returns. For these managers, communicating perceptions to investors could increase flows and assets under management. Our results suggest that investors do not fully see through managers' strategic communication choices, as they do in many rational expectations models (e.g., Stein, 1989, 2005). An alternative interpretation is that investors are aware of manager risk perceptions for all funds but choose to ignore them for funds with positive beta gaps and option-like equity exposures.

To better understand whether investors correctly anticipate hidden risk, we analyze flows during the first quarter of 2020, when public equity markets fell by almost 20% in response to the outbreak of the Covid-19 pandemic. Market drawdowns of this kind are useful for studying hidden risk because they are salient for investors and magnify the divergence between historical and manager-perceived betas. Consider a fund with a positive beta gap and assume investors are unaware of the manager's (valuable) private information. After a sharp market decline, such a fund is more likely to underperform expectations based on its historical beta, triggering abnormally large outflows. Indeed, at the onset of Covid, outflows from funds with a positive beta gap were strongly increasing in the size of the beta gap. This pattern, which contrasts sharply with behavior in normal times, suggests investors were caught off guard by the poor performance of these funds. Importantly, our tests control for realized CAPM alphas in 2020Q1, implying that investors were not simply reacting to weak performance but were instead revising their views on the risk of these funds.

Ex-post regret does not by itself imply investors could do better ex-ante. For instance,

if managers of positive-beta-gap funds keep their risk views private, it is not obvious how investors would identify such funds ahead of time. Nevertheless, we find that easily observable characteristics—most notably, a low historical CAPM beta—strongly predict a positive beta gap. Leverage, size, and past raw returns also matter. These patterns suggest that investors could improve their asset allocation decisions with information that is available to them ex-ante.

Overall, our findings indicate that hedge fund managers hold valuable risk information that is not reflected in past returns. Moreover, the link between flows and performance implies that investors do not always act on this private information, suggesting managers sometimes conceal their views. A back-of-envelope calculation implies that this agency distortion is fairly large in aggregate: according to our estimates, roughly 10-20% of flows into the hedge fund sector since 2015 may have been allocated differently if investors had access to managers' private risk assessments.

Our paper belongs to the broad literature on agency problems in asset management. Following Jensen and Meckling (1976) and Holmström (1979), the key underlying friction in much of this literature is moral hazard—investors cannot observe the manager's actions. Makarov and Plantin (2015) and Acharya et al. (2016) are recent examples that apply this idea specifically to models of manager risk-taking. A related theoretical literature studies the optimal contract between investors and fund managers in the face of such information frictions, including Bhattacharya and Pfleiderer (1985), Stein (2005), Panageas and Westerfield (2009), He and Xiong (2013) and Buffa et al. (2022).

Although agency problems in asset management are relatively well-understood in theory, they are challenging to measure empirically because researchers seldom observe the distinct information sets of managers and investors. Past studies have therefore relied on indirect approaches. For instance, Brown et al. (1996), Chevalier and Ellison (1997), Huang et al. (2011), and Han et al. (2021) show that the observed risk of mutual funds is positively correlated with variables that proxy for the personal incentives of managers. Similarly, Lo (2001), Agarwal and Naik (2004), and Jurek and Stafford (2015) argue that hedge funds

with positive realized CAPM alpha may simply be taking option-like risks that are hard to detect in normal times. Relative to previous work, the unique nature of Form PF allows us to more directly observe the agency problem because it contains assessments of fund risk that are made privately by hedge fund managers. Our key finding is that investors frequently behave as though they are unaware of this valuable information, and the resulting distortions to capital allocation are economically meaningful.

The idea that asset managers may seek to hide risk also features prominently in policy discussions (e.g., Rajan, 2005, Acharya et al., 2009). In a narrow sense, our results suggest these concerns may be warranted in the U.S. hedge fund sector. More generally, if hidden risk emerges even among hedge funds that largely invest in publicly traded securities, then agency problems are likely more severe in settings where risk is harder to observe from past returns—such as structured products (Coval et al., 2009), private equity and credit (Stafford, 2021), and real estate (Couts et al., 2020). This observation is especially pertinent given the rapid expansion of alternative asset classes over the past decade (Ivashina and Lerner, 2018, Greenwood et al., 2024, Begenau et al., 2025).

The remainder of the paper proceeds as follows. Section 2 provides background on Form PF and describes our analysis sample. Section 3 analyzes managers' private risk assessments and shows they provide valuable information that is not captured by past returns. Section 4 documents that investor flows respond to this information in funds with a negative beta gap, but not vice versa. This asymmetric response is consistent with strategic communication on the part of managers. Section 5 asks whether investors could improve their asset allocation decisions. Finally, Section 6 evaluates the aggregate size of this agency distortion and concludes.

2 Background and Data

In this section, we outline our data sources, describe our analysis samples, define key variables, and present basic summary statistics.

2.1 Form PF

Our fund-level data on hedge funds primarily comes from the Securities and Exchange Commission's (SEC) Form PF, which was adopted in 2011 to implement parts of Title IV of the Dodd-Frank Wall Street Reform and Consumer Protection Act of 2010. Investment advisers registered with the SEC that advise one or more private funds and have at least $150 million in gross assets under management in private funds must periodically file Form PF. The information contained in Form PF is confidential and is designed, in part, to assist regulators in their assessment of systemic risk in the U.S. financial system. Willful misstatements or omissions of material facts in a Form PF report are unlawful under section 207 of the Investment Advisers Act of 1940 and may result in the adviser's registration being revoked or criminal prosecution.

Hedge fund advisers file Form PF at least annually and report items such as gross and net asset values, returns, borrowings, and strategies. Large hedge fund advisers – those with at least $1.5 billion in gross assets managed in hedge funds – are required to file Form PF quarterly and additionally report more detailed information about their asset class exposures, risk metrics, asset and funding liquidity, and counterparties, among other items, for each of their large ("qualifying") hedge funds.[2] We keep only qualifying hedge funds in our sample because our analysis requires information provided by them, including their sensitivities to the equity market factor in Question 42, which we use to infer their managers' risk perceptions. As of Q3 2023, qualifying hedge funds managed 80 (86) percent of the net (gross) assets of all hedge funds filing Form PF.[3] We further retain only data between January 2013 and July 2023. We start the our analysis in 2013 because this was the first full year in which hedge funds were required to report.

[2]A qualifying hedge fund has a net asset value of at least $500 million as of the last day in any month in the fiscal quarter preceding the adviser's most recently completed fiscal quarter.

[3]See the SEC's Private Fund Statistics: https://www.sec.gov/divisions/investment/private-funds-statistics

2.2 Variable Construction and Analysis Sample

In this section, we define several key variables for our analysis. The unit of observation in our primary analysis sample is fund-quarter, as fund flows can only be computed at this frequency since net asset values (NAVs) are reported quarterly. Assets under management (AUM) is defined based on each fund's net asset value as reported in Question 9. Net quarterly flows into each fund are constructed using observed AUMs and quarterly net-of-fee returns. Following common practice in the literature, we handle outliers by trimming. We trim the top and bottom 1% of flows.

There are several different measures of leverage that one can extract from the PF data. We focus on gross-to-net asset value. In a corporate setting, this measure is akin to the ratio of assets to equity. Gross asset values are reported in Question 8 of Form PF. We trim values below 1 and the 99% tail of leverage.

Question 50 of Form PF asks for the fraction of each fund's capital base that can be withdrawn over one of the following fixed windows: 1 day or less, 2-7 days, 8-30 days, 31-90 days, 91-180 days, 181-365 days, and more than 365 days. We define the length l of each window using its midpoint. For example, we set the length of the 8-30 day window to be $l = 19$ days. The window exceeding 365 days is given a length of $l = 365$ days. The share liquidity of the fund is then a weighted average of these windows, with weights based on the fraction of capital $f(l)$ that can be withdrawn over that window.

Data on the fraction of each fund that is owned by insiders is taken from Form ADV, which is used by hedge managers to register with the SEC. When initially filing form ADV with the SEC, managers are required to provide a variety of information on each of their funds, including the fraction of each managed fund owned by insiders.[4] Moreover, Form ADV must be updated annually according to SEC regulation. For each fund-quarter observation in Form PF, we obtain the inside ownership share based on the ADV filling for that effective year.

Fund strategy is inferred based on Question 20, which asks for the percent of assets

[4]Inside ownership is taken from Question 14 of Form ADV, Part 1A, Schedule D, Section 7.B.(1).A.

managed in one of twenty two possible strategies, including a sector for "Other." For example, one possible strategy is "Equity, Market Neutral." We further collapse these granular strategies into one of the following nine high-level strategies: (i) equity; (ii) macro; (iii) relative value; (iv) event-driven; (v) credit; (vi) managed futures; (vii) investments in other funds; (viii) multi-strategy; and (ix) other. Funds are assigned to a particular strategy if at least 70% of their assets are in associated granular strategies. Funds that do not meet this criteria are considered multi-strategy.

Table 1a provides some basic summary statistics for our quarterly analysis sample. The table reflects only funds that report enough information for us to compute the manager's perceived exposure to equity markets. We describe the procedure by which we do so below in Section 3.1.1. Funds are also asked to report in Form PF their gross-of-fee and net-of-fee returns on a monthly basis. We use this monthly data to construct rolling 36-month estimates of beta and alpha with respect to the CRSP Value-Weighted index, requiring at least 24 months of valid return observations. Table 1b shows summary statistics for these monthly variables. Realized fees in Table 1b equal the difference between gross and net-of-fee returns.

3 Manager Risk Perceptions

In this section, we describe our approach to inferring a manager's perceptions about the risk of its funds using data from Form PF. Throughout the paper, we define risk in the classic CAPM sense, namely as the beta of the fund with respect to the aggregate stock market. We then compare perceived CAPM betas with the beta implied by historical returns, finding that this "beta gap" is economically meaningful and is not merely an artifact of estimation error. Moreover, we show that managers' perceptions of risk contain valuable information about future market exposure, offering predictive content beyond that of historical returns alone.

3.1 Basic Properties

3.1.1 Reporting in Form PF

We infer manager perceptions about risk based on their response to Question 42 in Form PF, which is stated as follows:

> For each of the market factors identified below, determine the effect of the specified changes on the reporting fund's portfolio and provide the results.

There are eight market factors that each fund is asked to consider: (i) equity prices; (ii) risk-free interest rates (parallel shifts in the yield curve); (iii) credit spreads; (iv) currency rates; (v) commodity prices; (vi) option-implied volatilities; (vii) default rates on asset-backed securities (ABS); and (viii) default rates on corporate bonds and credit default swaps (CDS). For each factor, the question specifies what it means by a change in the factor. For example, for equity prices, it says:

> A change in "equity prices" means that the prices of all equities move up or down by the specified amount, without regard to whether the equities are listed on any exchange or included in any index.

The fund is then asked to separately provide the effect on the long and short components of its portfolio, as a percent of NAV, if the market factor goes up or down by a pre-specified amount. For the equity-market factor, these amounts are: (i) an increase of 5%; (ii) an increase of 20%; (iii) a decrease of 5%; and (iv) a decrease of 20%.

Fund managers can answer in one of three ways. First, they can directly answer the question by reporting how the fund will perform under each factor scenario. Second, they can specify that the factor is not relevant for their portfolio and leave the question blank. Third, they can specify that the factor is relevant to the portfolio, but not tested as part of the fund's risk management, and again leave the question blank.

Table 2a shows response rates for each of the eight market factors. In the table, R-T stands for "relevant and tested" and corresponds to cases in which the fund manager provided a complete response; "R-NT" corresponds to cases in which the manager reports the

factor is relevant but not tested by their risk-management protocols; and "NR" corresponds to cases in which the manager reported the factor was not relevant for their portfolio. The unit of observation is fund-quarter, as funds are required to answer the question on Form PF quarterly.

A few facts stand out from Table 2a. First, the most frequently reported exposure by hedge funds is to the equity market factor, with 75% of respondents reporting it as relevant and 47% reporting it as both relevant and tested. The funds that report equities as relevant but not tested tend to have a higher proportion of their investments in non-listed equities. After equity market exposure, the two most relevant factors are currencies and rates.

For the remainder of the paper, we focus exclusively on equity-market exposure since it is the factor that is relevant and tested for the largest number of funds. The concept of equity-market risk we consider is CAPM beta, which Agarwal et al. (2018) show is the measure of risk that best explains fund flows into hedge funds. To compute CAPM beta from Form PF, we proceed as follows. Let $\beta(x)$ be a function that yields a manager's perceived beta given an aggregate market return of x. For $x \in \{-20\%, -5\%, 5\%, 20\%\}$, we set $\beta(x) = R(x)/x$, where $R(x)$ is how the manager expects the fund to perform given a market return of x.[5] For instance, if a manager reports that the fund's net long and short positions will fall by 10% when the market falls by 20%, then $R(-20\%) = -10\%$ and $\beta(-20\%) = 0.5$. If $x \notin \{-20\%, -5\%, 5\%, 20\%\}$, we use linear interpolation to determine beta when $|x| \leq 20$. For example, if $x = -10\%$, then $\beta(-10\%) = (1/3) * \beta(-20\%) + (2/3) * \beta(-5\%)$. We further assume that $\beta(x) = \beta(-20\%)$ for $x < -20\%$ and $\beta(x) = \beta(20\%)$ for $x > 20\%$.

In each month t within quarter $q(t)$, we use each fund's $\beta(x)$ function to define the manager's perceived beta based on the realized market return in t. Extending our notation from above, let $\beta_{i,q}(x)$ be the beta function reported by the manager of fund i at the end of quarter q. In a given month t, if the market return is $R_{m,t}$, then we define the perceived beta for i in month t as $\beta^R_{i,t} = \beta_{i,q(t)-1}(R_{m,t})$. That is, in quarter t, we use the beta function from the previous quarter and plug in as an input the realized market return in each month.

[5] We trim outliers of $\beta(x)$ each value of $x \in \{-20\%, -5\%, 5\%, 20\%\}$ at the 1% level.

This means $\beta_{i,t}^R$ is forward looking in the sense that it is based on beta function reported by the manager on Form PF from the previous quarter.

The remaining issue is how to compute $\beta_{i,t}^R$ for funds that report that the equity-market factor is either not relevant (NR) or relevant but not tested (R-NT). We think it is most natural to assume that $\beta_{i,t}^R = 0$ for those that say the factor is not relevant. For the remaining funds, we treat $\beta_{i,t}^R$ as missing.

Table 2b provides summary statistics on the reported beta function $\beta_{i,q}(x)$ across fund-quarter observations. In the table, we compute summary statistics only for the values of $\beta_{i,q}(x)$ that are directly reported in Form PF, so for $x \in \{-20\%, -5\%, 5\%, 20\%\}$. The average $\beta_{i,q}(x)$ in the sample is around 0.2, regardless of x. This is similar to the unconditional average CAPM beta for hedge funds of 0.37 reported by Barth et al. (2023). There is also substantial heterogeneity in betas: for example, at $x = -5\%$, the 5th percentile fund reports zero beta whereas the 95th percentile fund reports a beta of one. The fifth row of the table explores non-linearities in beta during market downturns by computing the ratio of $\beta_{i,q}(-20\%)/\beta_{i,q}(-5\%)$. A ratio of one suggests a linear relationship between beta and market returns, while a ratio greater than one indicates greater exposure during larger downturns compared to smaller ones. The data show that the median fund exhibits a linear beta, although there is again some variability across funds.

Figure 1 reports a variance decomposition of $\beta_{i,t}^R$. As a reminder, $\beta_{i,t}^R$ is based on the manager's reported beta function at the end of the previous quarter and the realized market return in month t. The variance decomposition in the figure reflects the adjusted R^2 from a regression of $\beta_{i,t}^R$ on one of several fixed effects. The plot shows that fund fixed effects explain over 80% of the variation in $\beta_{i,t}^R$, suggesting that perceptions of beta are relatively stable within each fund. Adviser fixed effects explain around 60% of the variation in $\beta_{i,t}^R$, consistent with the idea that advisers use the same risk management and assessment tools across all of the funds they oversee. Strategy appears to explain a modest amount of variation in perceived beta. Little variation in $\beta_{i,t}^R$ is absorbed by fixed effects based on deciles of each fund's leverage, as measured by gross-to-net asset value.

3.2 The Beta Gap

In this subsection, we compare the CAPM beta perceived by fund managers, β^R, to the one implied by historical fund returns, $\hat{\beta}$. This difference is defined as the beta gap, $g = \beta^R - \widehat{\beta}$. We then show that the beta gap varies widely across funds and is not simply a product of estimation error in $\hat{\beta}$.

3.2.1 Unconditional Estimates

To start, we estimate each fund i's CAPM beta, $\widehat{\beta}_i$, by regressing its full history of monthly excess returns on the excess return of the CRSP value-weighted (VW) index. Excess returns are computed relative to the one-month Treasury bill rate, and we use the fund's gross returns to make the estimated beta comparable to what funds report in Form PF. Using the full sample to estimate $\widehat{\beta}_i$ maximizes precision, but at the cost of ignoring potential time-series variation in observable market risk. In later sections, we allow for $\widehat{\beta}_i$ to vary through time using rolling regressions.

We then compare the estimated $\widehat{\beta}_i$ to each fund manager's average perceived beta, $\bar{\beta}_i^R$, which is also computed using each fund's full history in Form PF. The kernel density in Figure 2 visualizes the cross-fund distribution of the resulting beta gap, $g_i = \bar{\beta}_i^R - \widehat{\beta}_i$. Outliers of g_i are removed by trimming its 1% tails. Notably, the distribution is centered roughly around zero, with a mean of -0.06, suggesting that on average, fund managers' perceptions of risk align with historical returns. The figure also reveals significant dispersion in the beta gap, g_i: some managers perceive their fund's CAPM betas to be markedly different from that implied by its returns. The 10th percentile fund has a negative beta gap of -0.44, implying that the manager of this fund perceives lower risk than what its returns suggest. In contrast, the 90th percentile fund has a beta gap of 0.31, indicating the manager perceives higher risk than its returns suggest. The standard deviation of the beta gap is 0.33.

A simple way to understand the magnitude of the beta gap is to consider the problem of determining a fund's risk-adjusted return (alpha), assuming for simplicity that the market

is the only priced factor in the economy. From the manager's perspective, the fund's alpha equals $\alpha_i^R = \mu_i - \beta_i^R \mu^M$, where μ_i and μ^M are the average excess returns of the fund and market, respectively. Similarly, based on historical returns, the alpha of the fund equals $\widehat{\alpha}_i = \mu_i - \widehat{\beta}_i \mu^M$. The difference between the two alphas equals $\alpha_i^R - \widehat{\alpha}_i = -(\beta_i^R - \widehat{\beta}_i)\mu^M = -g_i \mu^M$. Intuitively, a positive beta gap g_i means the manager believes the fund's beta is relatively high and thus attributes more of the fund's average return to market exposure, not alpha. For a market risk premium of $\mu^M = 5\%$, a beta gap of $g_i = 0.3$ indicates the manager's perceived alpha is 1.5% per year lower than what is implied by historical returns. The realized market risk premium over our sample is 11.5% per year, implying that the gap between this manager's perceived alpha and the alpha estimated in historical returns is even larger in sample at nearly 3.5% per year.

While the heterogeneity in Figure 2 is sizable, it could simply be an artifact of estimation error in either β_i^R or $\widehat{\beta}_i$. To investigate this possibility more formally, we run the following pooled regression, using all fund-month observations:

$$\tilde{R}_{i,t+1} - \beta_{i,t+1}^R \tilde{R}_{M,t+1} = \theta_{0,i} + \theta_{1,i} \tilde{R}_{M,t+1} + \varepsilon_{i,t+1}, \tag{1}$$

where $\tilde{R}_{i,t+1}$ is fund i's realized excess return in month $t+1$, $\beta_{i,t+1}^R$ is the manager's reported beta in month $t+1$, and $\tilde{R}_{M,t+1}$ is the realized excess return of the CRSP VW index in month $t+1$. As discussed in Section 2.2, $\beta_{i,t+1}^R$ is the manager's estimate of beta based on its most recent response on Form PF and conditional on the realized market return in $t+1$. Consequently, the term $\tilde{R}_{i,t+1} - \beta_{i,t+1}^R \tilde{R}_{M,t+1}$ on the left-hand side of the regression is the manager's perceived idiosyncratic return for the fund. $\theta_{0,i}$ are fund fixed effects, and the slope coefficients $\theta_{1,i}$ are allowed to vary by fund. Under the null hypothesis that each fund manager's perceived beta coincides with its historical realized beta, the coefficients $\theta_{1,i}$ in the regression should be jointly zero. However, we can easily reject this null at conventional levels ($p = 0.00$), cutting against the idea that the dispersion in Figure 2 is driven solely by estimation error. Moreover, on a fund-by-fund basis, we reject that the unconditional

beta gap g_i is equal to zero for 52% percent of funds.

3.2.2 Time-series Persistence

Next, we consider a second, complementary test of whether the beta gap is driven by estimation error based on its persistence within each fund. Specifically, for each fund i, we construct at most three non-overlapping sample periods: $\tau_i = \{1, 2, 3\}$. Period $\tau_i = 1$ for fund i is the first 36-month period for which we observe monthly fund returns. Period $\tau_i = 2$ is the second non-overlapping period for which we observe returns, and so on. Because our data has at most 132 monthly observations for any fund, each fund has at most three 36-month non-overlapping sample periods. Within each period, we then construct the beta gap as the last perceived beta minus the estimated beta. Similar results obtain when using the average perceived beta during each period.

Figure 3 shows a binscatter plot of the estimated beta gap from period τ_i against its value from the previous period, $\tau_i - 1$. The plot clearly shows that firms with a high beta gap in period $\tau_i - 1$ tend to have a high beta gap in period τ_i. The slope of the linear regression line in the plot equals 0.56 and is statistically different from zero at conventional levels (s.e. $= 0.03$), with standard errors based on clustering at the fund level. Under the null hypothesis that the beta gap is purely driven by estimation error in CAPM betas, $\widehat{\beta}_i$, this slope coefficient should be zero, since estimation error should be uncorrelated across non-overlapping windows. Instead, Figure 3 indicates significant persistence in the beta gap within individual funds through time. To get a sense of economic magnitudes, suppose that the true beta gap is a permanent characteristic of each fund and the variance of the estimation error is constant over time. Under these assumptions, the persistence coefficient of 0.56 is equal to the fraction of variation in the observed beta gap that comes from variation in the true beta gap.[6]

[6]Formally, let g_i be the true beta gap for fund i and $g_i + \epsilon_{i,t}$ be the observed beta gap. Then the persistence coefficient is $\frac{Cov[g_i+\epsilon_{i,t}, g_i+\epsilon_{i,t-1}]}{Var[g_i+\epsilon_{i,t-1}]} = \frac{\sigma_g^2}{\sigma_g^2+\sigma_\epsilon^2}$.

3.3 Are Manager Perceptions Informative?

The evidence from the preceding subsection suggests meaningful differences between the risk implied by a fund's historical returns and the risk its manager perceives. Given these differences, it is natural to ask if managers have information about fund risk taking that is not embedded in historical returns. To explore this question, we now test whether a manager's perceived market beta $\beta^R_{i,t}$ helps predict the fund's equity market exposure, after controlling for the information in past realized returns. Specifically, we estimate the regression:

$$\tilde{R}_{i,t+1} = \gamma_0 + \gamma_1(\beta^R_{i,t+1}\tilde{R}_{M,t+1}) + \gamma_2(\widehat{\beta}_{i,t}\tilde{R}_{M,t+1}) + \varepsilon_{i,t+1}, \tag{2}$$

where $\tilde{R}_{i,t+1}$ is fund i's realized excess return in month $t+1$, $\beta^R_{i,t}$ is the manager's reported beta in month $t+1$ based on their most recent response to Form PF and the realized market return at $t+1$, $\widehat{\beta}_{i,t}$ is fund i's estimated CAPM beta based on realized returns between month $t-35$ and t, and $\tilde{R}_{M,t+1}$ is the realized excess return of the CRSP VW index in month $t+1$. The coefficient of interest in the regression is γ_1. If $\gamma_1 > 0$, the manager's perceptions of risk contain information that is not contained in past returns.

Figure 4a reports the results of regression (2) graphically. It plots a binscatter of $\tilde{R}_{i,t+1}$ against $\beta^R_{i,t+1}\tilde{R}_{M,t+1}$, after controlling for $\widehat{\beta}_{i,t}\tilde{R}_{M,t+1}$.[7] The figure shows that manager risk perceptions contain significant information about the fund's risk taking. The slope of the relationship is $\gamma_1 = 0.29$, and it is statistically significant at conventional levels with a standard error of 0.02 based on Driscoll and Kraay (1998) standard errors with fourteen lags.[8] For comparison, the sensitivity of realized excess returns to the fund's historical CAPM beta scaled by the realized market excess return is $\gamma_2 = 0.61$ (not depicted in the figure), and a perfect measure of the fund's equity risk exposure would recover a slope of 1. The upshot of Figure 4a is that manager perceptions offer economically meaningful

[7] Throughout the paper, we construct binned scatter plots following Cattaneo et al. (2024).

[8] The lag length was chosen according to Lazarus et al. (2018). Similar standard errors obtain when clustering by fund and time.

incremental information about fund risk taking beyond historical returns.

The parameters in regression (2) are identified from variation across funds and within funds over time. Figure 4b illustrates that both sources of variation lead to similar conclusions regarding the information content of managers' risk perceptions. This figure displays the estimated regression coefficients γ_1 and γ_2. The blue lines in the plot represent the baseline regression without additional controls, whereas the red lines depict the coefficients when including fund fixed effects, which isolate within-fund variation. The key observation is that the coefficients remain very similar across the two specifications, reinforcing the conclusion that fund managers posses information about risk that is not contained in historical returns.

4 How Do Investors Respond?

We now turn to the question of whether investors respond to the risk perceptions of fund managers. We start by laying out a conceptual framework to guide our empirical analysis. The main insight is that the risk adjustments made by investors can be understood by studying how their flows respond to past realizations of measured alpha (from returns) and the beta gap. This revealed-preference approach is related to those taken by Berk and van Binsbergen (2016), Barber et al. (2016), and Agarwal et al. (2018). The crucial difference is that we ask whether investor flows take into account the risk perceptions of fund managers, whereas past research has broadly studied the types of factors (e.g., value or momentum) investors consider when adjusting for risk.

4.1 Conceptual Framework

To fix ideas, consider a single fund and suppose that investors believe the fund's CAPM beta at each point in time is a weighted average of its measured beta (i.e, from historical returns) and the manager's perceived beta:

$$\beta_t^I = \omega \beta_t^R + (1 - \omega)\widehat{\beta}_t \tag{3}$$

The parameter ω controls how much weight investors put on the manager's beliefs when forming their own views. This weight could be zero for at least two reasons. First, investors may not believe the manager possesses useful information about the fund's risk, though our evidence from Section 3.3 suggests this view would not be correct on average. Second, the manager may not communicate its beliefs to investors, in which case ω will be zero (unless investors make strong inferences from the absence of communication). From the perspective of investors, the realized CAPM alpha of the fund is given by:

$$\alpha_t^I = \mu_t - \beta_t^I \mu_t^M,$$

where μ_t and μ_t^M are the average realized excess return of the fund and market, respectively. Similarly, define $\widehat{\alpha}_t = \mu_t - \widehat{\beta}_t \mu_t^M$ as the fund's measured alpha based on returns. From equation (3), the two alphas are related as follows:

$$\alpha_t^I = \widehat{\alpha}_t - \omega \underbrace{(\beta_t^R - \widehat{\beta}_t)}_{g_t} \mu_t^M \tag{4}$$

If the beta gap g is positive, investors perceive the fund's alpha to be lower than what is implied by historical returns alone. The amount they adjust the measured alpha when forming their own beliefs naturally depends on the beta gap (g_t), the realized market risk premium (μ_t^M), and the weight they put on the manager's beliefs (ω).

We further assume that investor flows into the fund are governed by the following equilibrium condition:

$$\begin{aligned} flow_{t+1} &= \theta \alpha_t^I + \epsilon_{t+1} \\ &= \theta \widehat{\alpha}_t - \theta \omega (\beta_t^R - \widehat{\beta}_t) \mu_t^M + \epsilon_{t+1} \end{aligned} \tag{5}$$

where ϵ_{t+1} is i.i.d. This equilibrium condition could be microfounded in a number of ways, for instance using the model of Berk and Green (2004). It is natural to assume that $\theta > 0$, so that flows are positive when investors believe the fund has positive CAPM alpha.

Another way to understand the equilibrium condition is as follows. The coefficient θ on $\widehat{\alpha}_t$ captures the response of flows to fund outperformance when it is measured using only realized returns. However, if investors also consider the risk perceptions of fund managers, their risk adjustment differs from the one used to compute $\widehat{\alpha}_t$. The difference is encapsulated by the term $(\beta_t^R - \widehat{\beta}_t)\mu_t^M$. When the beta gap is positive, it indicates that the manager perceives higher risk for the fund than what is implied by returns. Consequently, following positive market returns ($\mu_t^M > 0$), the manager's estimated alpha would be lower than the one implied by returns ($\widehat{\alpha}_t$). Therefore, if investors consider the manager's perceptions ($\omega > 0$), fund flows would be lower than those predicted solely by $\widehat{\alpha}_t$.

Next, consider a standard flow-performance regression that is augmented by the beta gap times the realized market risk premium:

$$flow_{t+1} = \kappa_0 + \kappa_1 \widehat{\alpha}_t + \kappa_2 (\beta_t^R - \widehat{\beta}_t)\mu_t^M + \nu_t \tag{6}$$

From equation (5), it is clear that the estimated regression coefficients will recover $\kappa_1 = \theta$ and $\kappa_2 = -\theta\omega$. In other words, we can use the regression to estimate $\omega = -\kappa_2/\kappa_1$. This simple logic is the basis for our subsequent empirical work. Note that it rests on two assumptions, namely that investors: (i) form beliefs according to (3); and (ii) the relationship between flows and perceived alpha is approximately linear, as in (5). The first assumption is arguably the stronger one, though it is reasonable to think that investors are unlikely to have information about a fund's risk that is not spanned by returns and the manager's information set. We consider an extension of this basic setup in which investors possess private information in the online appendix.

4.2 A Flow-Based Test: Baseline

Motivated by the conceptual framework laid out above, we now estimate the following flow-performance regression to study how, if at all, manager perceptions of risk influence

investor behavior:

$$flow_{i,q+1} = \theta_0 + \theta_1 \hat{\alpha}_{i,q} + \theta_2 \left(\beta^R_{i,q} - \hat{\beta}_{i,q}\right) \times \mu^M_{i,q} + \Gamma X_{i,q} + \varepsilon_{i,q+1}, \tag{7}$$

where $flow_{i,q+1}$ is the flow into fund i in quarter $q+1$, $\hat{\alpha}_{i,q}$ is the fund's realized CAPM alpha, measured from a 36-month rolling regression using monthly returns through the end of quarter q. $\beta^R_{i,q}$ is the manager's reported beta for the last month of quarter q. $\mu^M_{i,q}$ is the average market excess return measured over the same 36-month window as $\hat{\alpha}_{i,q}$ and $\hat{\beta}_{i,q}$.[9] $X_{i,q}$ is a vector of controls, including lagged quarterly flows and log AUM, along with liquidity and strategy-by-time fixed effects, consistent with the standard approach in the literature.[10] Standard errors in the regression are based on Driscoll and Kraay (1998) with eight lags. The choice of lag length is based on the recommendations of Lazarus et al. (2018).

Table 3 presents the results. In column 1, the coefficient on $\hat{\alpha}_{i,q}$ is $\theta_1 = 1.80$ and is strongly statistically significant (s.e. $= 0.24$). Since the $\hat{\alpha}_{i,q}$ is monthly and $flow_{i,q+1}$ is quarterly, this implies that outperformance of one percent translates into monthly flows of roughly $\theta_1/3 = 0.60\%$ of assets. These magnitudes are consistent with others in the literature (e.g., Fung et al., 2008, Getmansky, 2012, Barth et al., 2023).

The coefficient on $(\beta^R_{i,q} - \hat{\beta}_{i,q}) \times \mu_{m,q}$ is $\theta_2 = -0.66$. As predicted, the coefficient is negative, suggesting that on average investors do take manager risk perceptions into account. The coefficient is also statistically significant with a standard error of 0.26. As we will see below, this baseline regression masks significant heterogeneity across funds in θ_2. Investors in some funds respond strongly to manager risk perceptions, while investors in others do not respond at all.

The row corresponding to ω of Table 3 interprets the estimates of θ_1 and θ_2 through the lens of the framework laid out in Section 4.1. Recall that in the framework, the ratio $-\theta_2/\theta_1$

[9] The sample is restricted to funds with a minimum of 24 monthly returns over the 36-month window. The i subscript on $\mu^M_{i,q}$ reflects the fact that we measure the average market excess return only for months where the fund reports returns.

[10] See Sirri and Tufano (1998), Agarwal et al. (2018), and Barber et al. (2016) as just a few examples.

recovers the weight ω that investors put on managers' perceptions. The two coefficients imply that, on average, this weight is $\omega = 0.37$ with a standard error of 0.14. The standard error for the estimated ω is computed using the delta method.

The remaining columns of Table 3 probe the robustness of our baseline estimates in column (1) to the inclusion of a variety of controls. In column (2), we add controls for lagged size, the log of each fund's net asset value (NAV) last quarter, and lagged flows. The addition of these controls slightly lowers the magnitudes of θ_1 and θ_2, implying a similar ω to the specification in column (1).

Column (3) addresses the fact that flows into a fund are partly determined by its redemption policies (Liang et al., 2019). To account for share liquidity, we divide funds into deciles in each quarter based on the weighted average number of days needed for investors to redeem capital. The regression then includes fixed effects based on these decile assignments, meaning the coefficients are identified by comparing funds with similar share liquidity. The point estimates in the table show that controlling for liquidity does little to alter our main conclusions.

In column (4), we instead add strategy-by-time fixed effects, which absorb any time-series variation in fund flows that occurs at the strategy level. This means the coefficients of interest, θ_1 and θ_2, are identified from variation within a strategy in a given quarter. As discussed in Section 2.2, strategy definitions are based on Question 20 of Form PF and can be one of nine types, for instance equity or credit. The last column in Table 3 adds all fixed effects and controls and reassuringly finds similar results.

In sum, we find that for the average fund, investors flows appear to weakly respond to manager perceptions about risk. From the perspective of our model, we find that investors put a weight of roughly 0.4 on manager risk perceptions when adjusting fund performance for risk. Next, we explore whether this effect varies across different types of funds.

4.3 Heterogeneity

In this subsection, we document substantial heterogeneity in how investor flows respond to fund managers' perceptions of risk. We start by highlighting heterogeneity non-parametrically and then proceed to characterize the types of funds for which investors respond more strongly to manager perceptions. Our main finding is that investors flows are much less responsive to manager beliefs in funds with a positive beta gap, i.e., when the manager believes the fund is riskier than its returns suggest. Because manager incentives to communicate their beliefs are decreasing with the beta gap, these results suggest managers strategically reveal their beliefs about risk to attract flows. Consistent with this idea, investors appear to put more weight on manager beliefs when agency frictions are lower, namely when managers hold a larger stake in the fund.

4.3.1 Non-Parametric Evidence

We first take the following non-parametric approach. We estimate equation (7) separately for each fund f and compute the implied ω_f. Next, we sort funds into two groups based on whether their implied ω_f is above or below median. Finally, to improve precision, we re-estimate equation (7) separately for high- and low-ω funds, obtaining an estimate of ω for each group.

Figure 5 plots the resulting estimates of ω along with 95% confidence bands. The figure reveals that there is substantial variation in ω across funds, thus explaining the relatively low ω that we estimate in Table 3 when pooling across all funds. The estimated ω for high-ω funds equals 1.11 and has a standard error of 0.22. The fact that the point estimate is close to one implies that investor flows into these funds are very sensitive to manager perceptions of risk.

In contrast, low-ω funds have an estimated ω of 0.10. The standard error of this estimate equals 0.19, indicating that it is a relatively precisely estimated zero. The standard error is also such that we can easily reject the null that ω is equal across the two group at standard confidence levels. There are two possible reasons for why ω is close to zero for this group

of funds. The first is that investors are aware of manager risk perceptions but ignore them when risk-adjusting the returns of the fund and determining flows. The second is that managers of low-ω funds do not share their perceptions of risk with investors. In either case, investor flows would be less sensitive to the beta gap in equation (7), implying a low ω.

4.3.2 Positive vs. Negative Beta Gap Funds

In light of the substantial heterogeneity in ω documented in Figure 5, we now explore the types of funds for which ω is high or low. We focus specifically on variation in ω across funds with a positive or negative beta gap. The idea of splitting the sample based on the beta gap is simple: funds with a positive beta gap are those whose managers perceive higher risk than what is implied by the fund's returns. This in turn implies that the manager's perceived alpha is lower than the fund's measured alpha $\hat{\alpha}_{i,q}$, giving managers an incentive to withhold their view of the fund's risk from investors. The implication of this logic is that ω should be lower for funds with a positive beta gap.

Table 4 explores this hypothesis by estimating equation (7) for funds with a positive or negative beta gap. We define positive or negative beta gap funds using the same beta gap $g_{i,q} = \beta^R_{i,q} - \hat{\beta}_{i,q}$ that appears in equation (7), allowing for time-series variation in both measured and perceived beta. The first column examines funds with positive beta gaps. To reiterate, these are funds whose managers perceive more risk than implied by recent realized returns. The coefficient θ_2 on $(\beta^R_{i,q} - \hat{\beta}_{i,q}) \times \mu_{m,q}$ is close to zero and insignificant, implying that flows into these funds are not sensitive to the lagged beta gap. Unsurprisingly then, the implied weight ω that investors place on manager perceptions when risk-adjusting the returns of the fund is also close to zero. The second column of Table 4 shows that the estimates for positive beta-gap funds are robust to adding controls for size, lagged flows, and investor redemption restrictions, as well as strategy-by-time fixed effects.

Column (3) of Table 4 examines funds with negative beta gaps, i.e., funds whose managers perceive less risk than what is implied by recent returns. In contrast to positive beta

gap funds, the coefficient θ_2 for these funds is -1.38 and strongly statistically significant (s.e. $= 0.34$). To understand the intuition of the negative sign of θ_2, suppose that the average market return $\mu^M_{i,q}$ in equation (7) is positive and investors put a non-zero weight on manager perceptions of beta. For a negative beta-gap fund, the investor's perception of alpha will be then higher than the fund's measured alpha, $\hat{\alpha}_{i,q}$. Consequently, flows into the fund at $q+1$ should be higher than what is predicted from the measured alpha alone. Thus, θ_2 should be negative.

Combined with θ_1, the estimate of θ_2 in column (3) implies that investors put a weight ω of 0.69 (s.e. $= 0.23$) when risk-adjusting the returns of negative beta gap funds. Similar estimates obtain in column (4), when controlling for size, lagged flows, share liquidity, and strategy-by-time fixed effects. In all specifications, the magnitude of the implied ω for negative beta gap funds is substantially larger than that of positive beta-gap funds.[11]

While we cannot definitively pin down the microfoundations of these differences, they are consistent with strategic communication by fund managers about the risk of the fund. Managers with positive beta gaps perceive lower risk-adjusted performance than implied by their returns.[12] Thus, communicating their perceptions to investors could make these managers worse off by reducing flows and hence assets under management. Conversely, managers with negative beta gaps perceive higher risk-adjusted performance than implied by their returns. For these managers, communicating perceptions to investors could increase flows and assets under management. Under this interpretation, the differences we document in Table 4 also imply that investors are unable to fully see through managers strategic communication choices, as they would in some rational expectations models (e.g., Stein, 1989, Holmström, 1999)).

[11] Table A1 further shows that our conclusions about positive and negative beta-gap funds are robust to: (i) including fund and time fixed effects; (ii) clustering standard errors by fund and time, as opposed to using Driscoll and Kraay (1998); and (iii) defining the beta gap using rolling averages of perceived beta, $\beta^R_{i,q}$. See Section A.2 for more details.

[12] We confirm in Appendix A.1 that managers' perceptions of positive beta-gap funds are indeed informative about future fund risk, over and above the information embedded in past returns.

4.3.3 Non-Linear Betas

We next study funds with non-linear reported betas. More precisely, for each fund i, we define the following variables in each quarter q :

$$D_{i,q} = \frac{\beta_{i,q}(-20\%)}{\beta_{i,q}(-5\%)}$$
$$U_{i,q} = \frac{\beta_{i,q}(20\%)}{\beta_{i,q}(5\%)},$$

where $\beta_{i,q}(x)$ is the beta function that i reports in its Form PF response at the end of q (see Section 3.1.1). $D_{i,q}$ is simply the ratio of the fund manager's perceived beta conditional on a 20% market fall relative to its perceived beta under a 5% fall. $U_{i,t}$ is defined similarly, but for market booms.

We then define a variable, $O_{i,q}$, that equals one if either $D_{i,q}$ or $U_{i,q}$ exceeds one. $O_{i,q}$ indicates whether the manager perceives the fund's market exposure to increase with the absolute size of the market return. Such an exposure could arise, for instance, if the fund writes out-of-the-money put options or purchases out-of-the-money call options. The findings of Lo (2001) and Jurek and Stafford (2015) suggest that managers of funds with this type of nonlinear tail exposure have reduced incentives to communicate their risk perceptions to investors. This is because a fund that has exposure to market-wide tail events may appear to investors to earn CAPM alpha in normal times. Henceforth, we refer to funds for which $O_{i,q} = 1$ as nonlinear beta funds.

$O_{i,q}$ is a strong, positive predictor of the beta gap. A regression of an indicator $P_{i,q} = (\beta^R_{i,q} - \widehat{\beta}_{i,q}) > 0$ for whether fund i has a positive beta gap on $O_{i,q}$ indicates that nonlinear beta funds are 25.53 (s.e. $= 1.45$) percentage points more likely to have a positive beta gap. For context, recall that roughly half of the funds in our sample have a positive beta gap. Importantly, this result is not driven by a lack of large realized market returns during our sample, since the manager's perceived beta $\beta^R_{i,q}$ conditions on realized market returns. Instead, it implies that nonlinear beta funds are more likely to have positive beta gaps outside of market-wide tail events.

Columns (5)-(8) of Table 4 directly test whether investors in nonlinear beta funds respond to the beta gap by estimating equation (7) based on whether $O_{i,q} = 1$. Column (5) show estimates for nonlinear beta funds ($O_{i,q} = 1$). The coefficient θ_2 on the lagged beta gap interacted with the market return is positive and statistically insignificant. We obtain similar results in column (6), when adding controls and fixed effects for strategy and share liquidity. The implied ω in both specifications is around -0.2, and we cannot reject the null hypothesis that it is zero. Thus, in nonlinear beta funds, there is little evidence to suggest that investor flows account for manager perceptions of risk.

Column (7) instead estimates equation (7) for funds whose manager reports betas that are linear in market returns ($O_{i,q} = 0$). For these funds, θ_2 is negative and statistically significant. This remains true when we add controls and fixed effects in column (8). In contrast to nonlinear beta funds, the estimated ω in column (8) implies that investors assign a weight of 0.53 (s.e. = 0.11) on the manager's perceived beta when risk-adjusting returns. Together, the results in Columns (5)-(8) of Table 4 are consistent with the idea that managers of nonlinear beta funds strategically communicate their risk exposures with investors.

4.3.4 Inside Ownership

That managers have different beliefs than investors about the risk of the fund is a key principal-agent problem in asset management. Under standard agency theory (Jensen and Meckling, 1976), agency frictions of this sort can often be mitigated by skin in the game, i.e., having managers invest a meaningful fraction of their wealth in the fund. We now investigate this idea further by testing whether ω varies with the amount of inside ownership across funds. Inside ownership is measured based on Form ADV (see Section 2.2). To the extent that low values of ω reflect a strategic lack of communication by managers, we should observe higher levels of ω in funds that have high inside ownership.

Columns (1)-(2) of Table 5a run the baseline flow regression in equation (7) for funds with below-median inside ownership as of the end of quarter q. The first column shows

results when no other controls are added, and the second adds the controls, strategy-by-time fixed effects, and share liquidity fixed effects used in Table 3. Together, the two columns indicate that investor flows are not sensitive to manager perceptions of risk in funds with low inside ownership. The estimated loading θ_2 on the beta gap is near zero, implying an ω that is less than or equal to 0.15 and is not statistically significant.

Columns (5)-(6) rerun the regression for funds with above-median inside ownership. Focusing on column (6), which includes the whole suite of controls and fixed effects, the coefficient θ_2 on the beta gap is now negative and statistically significant. This suggests investors in these funds account for manager perceptions of risk. Indeed, the implied ω for this subset of funds is around 0.7 and statistically different from zero at conventional levels.

One concern with our results on inside ownership up to this point is that they may be somewhat mechanical. This is because the funds examined in these columns have higher inside ownership and insiders should respond to any private information they have about the fund's risk. Any easy way to deal with this potential issue is to compute flows into the fund that exclude insiders and re-run regression (7).[13] Columns (3)-(4) of Table 5a display the results using this modified definition of flows for funds characterized by low inside ownership. Conversely, columns (7)-(8) present the findings for funds with high inside ownership. Reassuringly, we continue to find similar results when excluding insider flows.

Overall, the results in Table 5a are consistent with the idea that high inside ownership aligns incentives between fund managers and investors. When inside ownership is high, the incentives for fund managers to strategically communicate their risk perceptions is lower. The fact that investors are more sensitive to manager risk perceptions in funds for which incentives to strategically communicate is low also further corroborates our interpretation of θ_2 and ω.

Table 5b further explores the interaction between the beta gap and inside ownership. To

[13]We do so by using the share of the fund owned by "Others" in Form PF, which includes inside ownership.

construct the table, we run regression (7) separately for different subsamples of funds based the intersection of the sign of their beta gap and whether they have above or below-median inside ownership. For instance, column (1) shows the regression, with no controls, for positive beta-gap funds that also have low inside ownership. In funds with low inside ownership, investors do not appear to consider manager risk perceptions when risk-adjusting returns, regardless of the sign of the beta gap. Within this set, there is some evidence in column (2) that negative beta-gap funds have a larger ω, but the estimates are quite noisy.

Columns (3) runs the regression for positive beta-gap funds with high inside ownership. Interestingly, we find that the implied ω is still rather low and statistically insignificant, suggesting that the inside ownership does not outweigh the incentives for managers to strategically communicate with investors in these funds. Perhaps unsurprisingly, investors are the most sensitive to manager perceptions of risk for funds with high inside ownership and a negative beta gap (column 4). For these funds, the incentives to strategically communicate with investors are the lowest.

4.4 Fees

Fees offer a complementary way to test whether investors account for manager risk perceptions when risk-adjusting returns. To see why, consider a simple scenario in which two funds, i and j, that launched at the same time and have the same measured gross beta ($\widehat{\beta}_i = \widehat{\beta}_j = \widehat{\beta}$) and alpha ($\widehat{\alpha}_i = \widehat{\alpha}_j = \widehat{\alpha}$). Suppose the manager-perceived betas of the two funds are such that $\beta_i^R > \widehat{\beta} > \beta_j^R$. This means i has a positive beta gap and j has a negative one. So long as investors in both funds put non-negative weights on manager risk perceptions, equation (4) implies that investors in fund i will perceive its alpha to be lower than in fund j. Consequently, they should only be willing to pay lower fees for fund i.

To operationalize this logic in our data, we first compute the average realized fee, fee_i, for each fund i based on the difference between its monthly gross-of-fee and net-of-fee returns. We also compute the unconditional beta gap $g_i = \bar{\beta}_i^R - \widehat{\beta}_i$ using each fund's full history of returns, as in Section 3.2.1. We then run the following cross-sectional regression

for different subsamples of funds:

$$fee_i = \lambda_0 + \lambda_1 \widehat{\alpha}_i + \lambda_2 (\bar{\beta}_i^R - \widehat{\beta}_i) \mu_i^m + \varepsilon_i, \tag{8}$$

where $\widehat{\alpha}_i$ is each fund's unconditional measured alpha and μ_i^m is the average excess return of the CRSP VW index over the fund's lifetime. Standard errors are clustered by adviser. All regressions include fixed effects for strategy, deciles based on each fund's average share liquidity, and deciles based on each fund's average NAV. Funds must have at least one year of data to be included in the analysis. Under the null hypothesis that investors account for manager risk perceptions, we should find $\lambda_2 < 0$ for all funds, regardless of the sign of their beta gap. If not, this is an indication that investors do not account for manager risk perceptions when determining fund fees.

Column (1) of Table 6 shows estimates of equation (8) for funds with a positive average beta gap. Unsurprisingly, the estimated λ_1 is positive and statistically significant, suggesting average fees are increasing in measured alpha. However, while λ_2 is negative, it is not statistically different from zero, implying that fees in positive beta-gap funds are not sensitive to the size of the beta gap. In contrast, column (2) shows that λ_2 is negative and statistically significant for funds with a negative beta gap, though the standard errors are such that we cannot reject the null that λ_2 is equal across the two subsamples.

In columns (3)-(4), we run the test for nonlinear beta funds, which we define unconditionally based on whether i's manager reports a non-linear beta $O_{i,q}$ for at least half of its sample. Column (3) shows that fees are not sensitive to the beta gap within funds with a nonlinear beta, whereas column (4) shows that they are for funds with a linear beta. Overall, our analysis of fees reinforces our previous findings that flows into positive beta-gap and nonlinear beta funds are not sensitive to manager perceptions of risk (Section 4.3.2). These results are consistent with the idea that managers in these funds strategically avoid communicating their perceptions of risk to investors.

5 Could Investors do Better?

The previous section provided evidence that investors in positive beta-gap funds do not take into account managers' risk perceptions. In this section, we first show that investors tend to reallocate away from positive beta-gap funds after extreme market returns, suggesting that investors regret their allocations to these funds ex post. Given that our evidence in Section 4 is consistent with the managers of these funds not sharing their perceptions with investors, it is not obvious that investors could have done better ex ante. We then show that observable fund characteristics correlate with the beta gap. Taken together, these two results suggest that investors could improve their ex ante asset allocations using available information.

5.1 Are investors surprised by fund performance?

When managers perceive their risk to be higher than what is implied by returns, the benefits to strategic communication are clear. If investors risk-adjust returns based solely on historical returns, they may overestimate the fund's true alpha. This perceived outperformance can attract more flows to the fund and potentially increase the fees charged by the managers.

The downside of this communication strategy is that investors may eventually recognize the actual risk levels of the fund and choose to withdraw their capital. This type of investor updating is most likely during periods of high market volatility, such as the onset of the Covid-19 pandemic in the first quarter of 2020. During this period, the CRSP VW index fell by over 20%, and the VIX index, a measure of stock market volatility, exceeded 50%. We therefore use this event to study investor updating.

The logic of our test can be readily understood through a simple example. Consider a fund with a positive beta gap and assume that entering the pandemic investors were either not aware of the manager's risk perceptions or chose to ignore them. When the stock market subsequently fell sharply in 2020Q1, the fund's performance was likely worse than investors expected. Consequently, flows out of the fund were likely abnormally high relative to other

funds. This argument of course assumes the manager has relevant information about the fund's risk, which we confirmed in Section 3.3.

Testing this hypothesis empirically requires a measure of abnormal fund flows. Following the literature (e.g., Guercio and Tkac, 2008), we first construct several measures using variants of the following flow-performance regression:

$$flow_{i,q+1} = \lambda_{s,l,h} + \sum_{s=0}^{3} c_s^r R_{i,q-s} + \sum_{s=0}^{3} c_s^f flow_{i,q-s} + \varepsilon_{i,q+1}, \tag{9}$$

where $R_{i,q}$ is the return of fund i in quarter q. Our baseline specification pools all funds together and includes a strategy-by-liquidity-by-size fixed effect. This fixed effect is constructed by first sorting funds into deciles based on their share liquidity in each quarter. We also sort funds into deciles based on their NAV. The fixed effect $\lambda_{s,l,h}$ is then based on the interaction of strategy (s) and indicators for the liquidity (l) and size (h) decile assignments.

In all cases, we estimate the flow-performance regression using data only through 2019Q4. This ensures the estimated parameters only reflect investor behavior before the pandemic. Using the fitted coefficients, we then construct a predicted flow $\widehat{flow}_{i,2020Q1}$ for each fund in 2020Q1. The abnormal flow in the quarter is then defined naturally as $\varepsilon_{i,2020Q1} = flow_{i,2020Q1} - \widehat{flow}_{i,2020Q1}$.

Armed with a measure of abnormal flows, we then run the following regression:

$$\varepsilon_{i,2020Q1} = a + b_1 1(g_i > 0) + b_2 g_i + b_3 1(g_i > 0) g_i + \nu_i \tag{10}$$

where $g_i = \bar{\beta}_{i,R} - \widehat{\beta}_i$ is the beta gap of fund i, also measured using all data through 2019Q4.[14] By the logic outlined above, we should observe $b_1 < 0$ if investors were on average surprised by the poor performance of positive beta-gap funds during 2020Q1. Moreover, the size of the surprise for funds with a positive beta gap should be increasing in the beta gap, meaning

[14] Our approach to testing equation (10) is essentially a two-step procedure, where abnormal flows are estimated in the first step. In Appendix A.3, we show that a one-step version of the test yields very similar results.

$b_3 < 0$.

Column (1) of Table 7 presents regression estimates of equation (10). The coefficient on b_1 is indeed negative, though is measured imprecisely. More strikingly, $b_3 = -14.92$ and is statistically significant at conventional levels (s.e. $= 4.19$), indicating that outflows from positive beta-gap funds are increasing in the size of the beta gap. The magnitude of this effect is large. Within the set of positive beta-gap funds, the 10th and 90th percentiles of the beta gap are 0.02 and 0.56, respectively. The estimated b_3 therefore implies that the 90th percentile fund experienced 8.02 percentage points more outflows than the 10th percentile fund.

Interestingly, the estimated b_2 is positive and statistically significant, although it is smaller in absolute magnitude than b_3 by a factor of two. This positive coefficient implies that, within the group of funds with a negative beta gap, funds with a more negative gap faced greater outflows. One possible explanation for this observation is that funds with a more negative beta gap experienced steeper declines during the market crash in the first quarter of 2020. This relatively poor performance then prompted investors to reassess the fund's beta, shifting their view away from the manager's perceived beta towards the higher beta implied by historical returns, and thus leading to larger outflows.

The remaining columns of Table 7 probe the sensitivity of the results in column (1) to different measures of abnormal flows, $\varepsilon_{i,2020Q1}$. In column (2), we allow the flow-performance relationship in equation (9) to vary by fund strategy, share liquidity, size, and the sign of the beta gap. We specifically do so by including only one lag of returns ($R_{i,q}$) and interacting it with the fixed effect $\lambda_{s,h,l}$ and the sign of the beta gap. Column (3) goes one step further by allowing the constant and the coefficient on lagged returns to instead vary by manager. In column (4), we use the same abnormal flows from column (3) and include a fixed effect for strategy and deciles of share liquidity in regression (10). Column (5) instead uses abnormal flows based on fund-by-fund regressions and uses no fixed effect in regression (10). The estimated b_2 is somewhat lower in this case compared to the other specifications, but that is likely driven by the fact that fund-level estimates of equation (9) yield noisier estimates

of abnormal flows.

5.2 Is the beta gap predictable?

In the preceding subsection, we show that investors appear to be negatively surprised by the performance of positive beta gap funds in extreme market downturns. This suggests that investors ex post regret their allocations to these funds. However, given that our evidence in Section 4 is consistent with the managers of these funds not sharing their perceptions with investors, it is not obvious that investors could have done better ex ante.

We now ask whether information available to investors can predict the beta gap. To guide our analysis of the factors predicting the beta gap, we first employ fixed effects regressions to identify its principal dimensions of variation. The results are summarized in Table 8a, which presents three key statistics from a regression of the beta gap $\beta_{i,t}^{R} - \hat{\beta}_{i,t}$ from Eq. (7) on different fixed effects. These statistics include: (i) the adjusted-R^2; (ii) the p-value from testing the null that the fixed effects are zero; and (iii) the p-value from testing the null that the fixed effects are equal. One potential issue with this variance decomposition is that the beta gap is mechanically persistent, since $\hat{\beta}_{i,q}$ is estimated using rolling regressions. In turn, the size of fixed effects may be artificially inflated. To avoid this issue, we use the same sample from Section 3.2.2, which samples each fund every 36 months to ensure that $\hat{\beta}_{i,t}$ is estimated using non-overlapping windows.

The first row of Table 8a reports the results for fund fixed effects. The adjusted R^2 in this regression is 54%, indicating that a large part of the beta gap can be thought of as a fund-level trait. This conclusion is consistent with Figure 3, which shows that the beta gap is persistent within each fund. The second column tests the hypothesis that all the estimated fund fixed effects are zero and finds that it is strongly rejected. The third column tests that the estimated fund fixed effects are equal to one another and finds that it is also strongly rejected. These rejections are not surprising given the large R^2 in column (1).

The second row shows results for adviser effects. Here, we find that 40% of the variation in the beta gap is explained by the identity of each fund's adviser, indicating that some

advisers consistently have a positive or negative beta gap across all of their funds and through time. The second and third columns show we reject the null of zero and equal adviser fixed effects, respectively.

Interestingly, the third row shows that variation in the beta gap is largely unrelated to the fund's strategy. The adjusted R^2 in this case is 2%. That said, there is enough variation in the beta gap across strategy to reject the null of zero and equal strategy effects. Finally, the last row in Table 8a asks whether there is common time-series variation in the beta gap across funds. The low adjusted-R^2 suggests that the answer is not much.

Next, we look for fund characteristics that are plausibly observable to investors and also predict the sign of the beta gap. We estimate regressions of the following form:

$$1(\beta_{i,t}^R - \hat{\beta}_{i,t} > 0) = a + bX_{i,t} + \varepsilon_{i,t}, \tag{11}$$

where $X_{i,t}$ is one of several different characteristics for fund i. The sample for the regression is the same as in Table 8a, which as a reminder contains non-overlapping 36-month windows for each fund. All standard errors are clustered within fund. For continuous variables, we standardize $X_{i,t}$ to have a mean of zero and standard deviation of one. Indicator variables are not standardized in this fashion. We scale the outcome variable by 100 to make the regression coefficients more interpretable. Table 8b contains the results.

The first characteristic we consider in column (1) is the beta implied by the fund's returns. To minimize any mechanical relationship between this beta and the sign of the beta gap, we use lagged values of measured beta. The table shows that the level of measured beta is a strong and negative predictor of the sign of the beta gap. A one standard deviation decrease in measured beta corresponds to roughly a -10 percentage point (s.e. $= 1$) increase in the likelihood that a fund has a positive beta gap. Recall that roughly half of funds have a positive beta gap. Even if some of the relationship in column (1) is mechanical, the broader point is that investors possess some information that could allow them to discern between positive and negative beta-gap funds ex-ante.

Columns (2) and (3) explore whether a fund's balance sheet structure offers predictive power for the sign of the beta gap. Column (2) focuses on balance sheet leverage, as measured by the ratio of gross and net assets. The table shows that funds with high leverage, defined as being in the top tercile of leverage, are 6 (s.e. = 2) percentage points more likely to have a positive beta gap. Column (3) shows that a one standard deviation increase in derivatives usage, defined as ratio of gross notional in derivatives to NAV, corresponds to a roughly 2 (s.e. = 1) percentage point increase in the likelihood of having a positive beta gap.[15]

In column (4), we explore whether fund size is associated with the beta gap. Specifically, in each month, we sort funds into terciles based on their last available NAV. We then define an indicator variable for funds in the top tercile of size. Column (4) shows that large funds are 6 percentage points (s.e. = 2) more likely to have a positive beta gap.

The last fund characteristic we consider is past performance. Past returns are defined using the previous 36-months of gross returns and funds are considered to have high returns if they are in the top tercile. The table shows that high past returns are a statistically significant and positive predictor of the sign of the beta gap. In terms of magnitude, high performing funds are 9 percentage points more likely to have a positive beta gap.

Overall, the results in Table 8b show that the sign of the beta gap is predictable by a number of observable fund characteristics, most notably the CAPM beta based on historical returns alone. In Table A3, we draw similar conclusions when predicting the level of the beta gap, not just its sign. Given that manager perceptions of risk contain valuable information (Section 3.3) and investors appear to ex post regret their allocations to positive beta gap funds, these findings suggest that investors in hedge funds could improve their asset allocations with available information.

[15]For both derivative usage and balance sheet leverage, we use the last available observed value for each fund-month observation. Gross notional of derivatives is based on Question 44 of Form PF. Outliers of derivative usage are trimmed at their 1% tails.

6 Implications and Conclusion

6.1 Counterfactual Flows

Given our results, it is natural to ask how fund flows might differ if investors accounted for manager beliefs when risk-adjusting returns. In the language of the model, we seek to compute counterfactual flows—and thus fund sizes—under different weights ω. Recall from Section 4.2 that we estimate ω via the following panel regression:

$$flow_{i,q+1} = \kappa_1 \widehat{\alpha}_{i,q} + \kappa_2 \left(\beta^R_{i,q} - \hat{\beta}_{i,q}\right) \times \mu^M_{i,q} + \Gamma X_{i,q} + \varepsilon_{i,q+1}, \tag{12}$$

where $\kappa_1 = \theta$ recovers how sensitive investor flows are to their own beliefs about fund alpha and $\kappa_2 = -\theta\omega$. Next, suppose that instead of the ω embedded in κ_2, investors put a weight ω_c on manager beliefs. Holding all other parameters fixed, including the idiosyncratic shock to flows $\varepsilon_{i,q+1}$, flows into the fund would then be given by:

$$\begin{aligned} flows^c_{i,q+1} &= \kappa_1 \widehat{\alpha}_{i,q} - \theta\omega_c \left(\beta^R_{i,q} - \hat{\beta}_{i,q}\right) \times \mu^M_{i,q} + \Gamma X_{i,q} + \varepsilon_{i,q+1} \\ &= flow_{i,q+1} - (\theta\omega_c + \kappa_2) \left(\beta^R_{i,q} - \hat{\beta}_{i,q}\right) \times \mu^M_{i,q}. \end{aligned} \tag{13}$$

To understand the intuition of this relationship, consider positive gap funds. As we showed in Section 4.3.2, roughly speaking, these funds have an $\omega = 0$, or equivalently, $\kappa_2 = 0$. This means that, in reality, investors put zero weight on manager beliefs. Now suppose we want to run the counterfactual in which $\omega_c = 1$. In this case, for positive beta-gap funds, investors in the counterfactual will perceive the alpha of the fund to be lower than they do in the data. Consequently, counterfactual flows in equation (13) will be lower than actual flows by an amount $\theta \left(\beta^R_{i,q} - \hat{\beta}_{i,q}\right) \times \mu^M_{i,q}$.

Equation (13) therefore suggests a simple two-step procedure for computing counterfactual flows for different ω_c. First, estimate (12) for any subset of funds and retain the estimated $\kappa_1 = \theta$ and κ_2. Second, plug these estimates and a desired counterfactual ω_c into equation (13) to compute counterfactual flows. From here, it is straightforward to

construct a time-series of counterfactual fund NAVs by combining observed fund returns and counterfactual flows $flows^c_{i,q+1}$.

Figure 6 displays a time-series of aggregate hedge fund flows under a few different counterfactuals. For comparison, the solid blue shows actual aggregate flows. All analysis for the plot is indexed to January 2015, as this is the first year for which we have sufficient data to estimate (12). The dotted blue line considers a counterfactual in which we set $\omega_c = 1$ for funds with positive beta gaps. In the language of the model from Section 4.1, this means we assume investors in these funds put full weight on the manager's perceived beta, which is higher than measured beta. Flows into negative beta-gap funds are held at their true values. In this counterfactual, total flows into the hedge fund sector would have been reduced by about 8 percentage points or roughly $35 billion in dollar terms.

The green dotted line in the plot instead holds fixed the actual flows into positive beta-gap funds and sets $\omega_c = 0$ for negative beta-gap funds. This counterfactual therefore reflects a scenario where investors in negative beta-gap funds ignore manager perceptions of risk when risk-adjusting returns, since perceived risk is lower than what is implied by returns. Under this counterfactual, flows into the hedge fund sector would have been reduced by over 10 percentage points.

Finally, the maroon dotted line combines both counterfactuals. That is, we set $\omega_c = 1$ for positive beta-gap funds and $\omega_c = 0$ for negative beta-gap funds. κ_1 is also allowed to differ across both groups and is based on subsample estimates of equation (12). The idea behind this counterfactual is to simulate how the industry would have evolved if investors were both fully aware of manager risk perceptions and skeptical when forming their beliefs about each fund's risk. By skeptical, we mean that investors assume the fund's risk is given the maximum of $\beta^R_{i,q}$ and $\hat{\beta}_{i,q}$. This means investors put full weight on each fund's measured beta for negative beta-gap funds and full weight on the manager's perceptions for positive beta-gap funds. In this case, flows into the hedge fund sector from January 2015 through June 2023 would have been reduced by roughly 20 percentage points ($84 billion).

This exercise is only suggestive for several reasons. First, our regressions are identified

in part from cross sectional variation and the usual concerns about extrapolating to the aggregate from the cross section apply. These concerns are particularly stark in this case, as we are assuming in our counterfactual that if a particular fund does not attract flows, those flows do not enter the hedge fund sector at all. An alternative interpretation of our calculation is that approximately 10-20% of flows would have been allocated differently across funds, with a smaller effect on the aggregate assets of the hedge fund sector. Second, we do not know what the correct level of ω is; it is likely to vary across funds and over time. Thus, our counterfactuals do not clearly correspond to any notion of the first best. Third, equation (12) is not a microfounded structural equation and is unlikely to the exactly capture counterfactual flows.

6.2 Conclusion

In this paper, we use novel data from SEC form PF to study how hedge fund managers perceive the risk of their funds. We show that for many funds, managers perceptions differ meaningfully from risk measured from historical returns. Consistent with strategic communication by fund managers, investor flows appear to take manager risk perceptions into account only when doing so increases the fund's implied outperformance. We provide evidence that investors update about fund risk after extreme market downturns and could have used available information to update earlier.

References

Acharya, Viral, Marco Pagano, and Paolo Volpin, "Seeking alpha: Excess risk taking and competition for managerial talent," *The Review of Financial Studies*, 2016, *29* (10), 2565–2599.

_ , Thomas Philippon, Matthew Richardson, and Nouriel Roubini, "The financial crisis of 2007-2009: Causes and remedies," *Restoring financial stability: How to repair a failed system*, 2009, pp. 1–56.

Agarwal, Vikas and Narayan Y Naik, "Risks and portfolio decisions involving hedge funds," *The Review of Financial Studies*, 2004, *17* (1), 63–98.

_ , T. Clifton Green, and Honglin Ren, "Alpha or beta in the eye of the beholder: What drives hedge fund flows?," *Journal of Financial Economics*, 2018, *127* (3), 417–434.

Barber, Brad M., Xing Huang, and Terrance Odean, "Which Factors Matter to Investors? Evidence from Mutual Fund Flows," *The Review of Financial Studies*, 06 2016, *29* (10), 2600–2642.

Barth, Daniel, Juha Joenväärä, Mikko Kauppila, and Russ Wermers, "A Bias-Free Assessment of the Hedge Fund Industry: A New Evaluation of Total Assets, Alphas, and the Flow-Performance Relation," Working Paper 2023.

Begenau, Juliane, Pauline Liang, and Emil N. Siriwardane, "The Rise of Alternatives," *Available at SSRN 4940886*, 2025.

Berk, Jonathan B. and Jules H. van Binsbergen, "Assessing asset pricing models using revealed preference," *Journal of Financial Economics*, 2016, *119* (1), 1–23.

_ and Richard C. Green, "Mutual Fund Flows and Performance in Rational Markets," *Journal of Political Economy*, 2004, *112* (6), 1269–1295.

Bhattacharya, Sudipto and Paul Pfleiderer, "Delegated portfolio management," *Journal of Economic Theory*, 1985, *36* (1), 1–25.

Brown, Keith C, W Van Harlow, and Laura T Starks, "Of tournaments and temptations: An analysis of managerial incentives in the mutual fund industry," *The Journal of Finance*, 1996, *51* (1), 85–110.

Buffa, Andrea M, Dimitri Vayanos, and Paul Woolley, "Asset management contracts and equilibrium prices," *Journal of Political Economy*, 2022, *130* (12), 3146–3201.

Cattaneo, Matias D., Richard K. Crump, Max H. Farrell, and Yingjie Feng, "On Binscatter," *American Economic Review*, May 2024, *114* (5), 1488–1514.

Chevalier, Judith and Glenn Ellison, "Risk taking by mutual funds as a response to incentives," *Journal of political economy*, 1997, *105* (6), 1167–1200.

Couts, Spencer J., Andrei S. Goncalves, and Andrea Rossi, "Unsmoothing Returns of Illiquid Funds," *SSRN Working paper*, 2020.

Coval, Joshua, Jakub Jurek, and Erik Stafford, "The Economics of Structured Finance," *Journal of Economic Perspectives*, March 2009, *23* (1), 3–25.

Driscoll, John C. and Aart C. Kraay, "Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data," *The Review of Economics and Statistics*, 1998, *80* (4), 549–560.

Fung, William, David A Hsieh, Narayan Y Naik, and Tarun Ramadorai, "Hedge funds: Performance, risk, and capital formation," *The Journal of Finance*, 2008, *63* (4), 1777–1803.

Getmansky, Mila, "The life cycle of hedge funds: Fund flows, size, competition, and performance," *The Quarterly Journal of Finance*, 2012, *2* (01), 1250003.

Greenwood, Robin, David S. Scharfstein, and Robert Ialenti, "The Evolution of Financial Services in the United States," HBS Working Paper, Harvard Business School 2024.

Guercio, Diane Del and Paula A. Tkac, "Star Power: The Effect of Monrningstar Ratings on Mutual Fund Flow," *Journal of Financial and Quantitative Analysis*, 2008, *43* (4), 907–936.

Han, Xiao, Nikolai L Roussanov, and Hongxun Ruan, "Mutual fund risk shifting and risk anomalies," *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper*, 2021.

He, Zhiguo and Wei Xiong, "Delegated asset management, investment mandates, and capital immobility," *Journal of Financial Economics*, 2013, *107* (2), 239–258.

Holmström, Bengt, "Moral hazard and observability," *The Bell Journal of Economics*, 1979, pp. 74–91.

_ , "Managerial incentive problems: A dynamic perspective," *The Review of Economic studies*, 1999, *66* (1), 169–182.

Huang, Jennifer, Clemens Sialm, and Hanjiang Zhang, "Risk shifting and mutual fund performance," *The Review of Financial Studies*, 2011, *24* (8), 2575–2616.

Ivashina, Victoria and Josh Lerner, "Looking for Alternatives: Pension Investments Around the World, 2008 to 2017," 2018. Working Paper.

Jensen, Michael C. and William H. Meckling, "Theory of the firm: Managerial behavior, agency costs and ownership structure," *Journal of Financial Economics*, 1976, *3* (4), 305–360.

Jurek, Jakub W. and Erik Stafford, "The Cost of Capital for Alternative Investments," *The Journal of Finance*, 2015, *70* (5), 2185–2226.

Lazarus, Eben, Daniel J. Lewis, James H. Stock, and Mark W. Watson, "HAR Inference: Recommendations for Practice," *Journal of Business & Economic Statistics*, 2018, *36* (4), 541–574.

Liang, Bing, Christopher Schwarz, Mila Getmansky Sherman, and Russ Wermers, "Share restrictions and investor flows in the hedge fund industry," *Available at SSRN 2692598*, 2019.

Lo, Andrew W., "Risk Management for Hedge Funds: Introduction and Overview," *Financial Analysts Journal*, 2001, *57* (6), 16–33.

Makarov, Igor and Guillaume Plantin, "Rewarding trading skills without inducing gambling," *The Journal of Finance*, 2015, *70* (3), 925–962.

Milgrom, Paul R., "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics*, 1981, *12* (2), 380–391.

Panageas, Stavros and Mark M Westerfield, "High-water marks: High risk appetites? Convex compensation, long horizons, and portfolio choice," *The Journal of Finance*, 2009, *64* (1), 1–36.

Rajan, Raghuram, "Has Financial Development Made the World Riskier?," in "Proceedings of the Jackson Hole Conference organized by the Federal Reserve Bank of Kansas City" 2005.

Sirri, Erik R. and Peter Tufano, "Costly Search and Mutual Fund Flows," *The Journal of Finance*, 1998, *53* (5), 1589–1622.

Stafford, Erik, "Replicating Private Equity with Value Investing, Homemade Leverage, and Hold-to-Maturity Accounting," *The Review of Financial Studies*, 02 2021, *35* (1), 299–342.

Stein, Jeremy C., "Efficient Capital Markets, Inefficient Firms: A Model of Myopic Corporate Behavior," *Quarterly Journal of Economics*, 1989, *104* (Dec.), 655–669.

_ , "Why are Most Funds Open-End? Competition and the Limits of Arbitrage*," *The Quarterly Journal of Economics*, 02 2005, *120* (1), 247–272.

Table 1: Summary Statistics for Analysis Sample

	Mean	Stdev	p10	p25	p50	p75	p90	N
AUM ($ billion)	2.11	3.66	0.22	0.56	1.04	2.15	4.69	31,169
Inside Ownership (%)	13.68	24.82	0.00	0.00	3.00	13.00	41.00	29,983
Leverage	1.71	1.34	1.00	1.03	1.26	1.73	2.86	30,398
Share Restrictions (days)	141	125	4	19	95	247	365	31,166
Net Flows (%)	-2.45	10.96	-13.26	-5.45	-0.87	1.09	6.42	30,550
Number of advisers	620							
Number of funds	2,111							

(a) Quarterly sample

	Mean	Stdev	p10	p25	p50	p75	p90	N
Gross return (%)	0.66	4.44	-3.03	-0.76	0.49	1.91	4.32	91,757
Net return (%)	0.47	4.09	-2.99	-0.80	0.38	1.68	3.95	91,670
Fee (%)	0.19	0.85	0.00	0.00	0.08	0.19	0.47	91,670
Measured CAPM β	0.33	0.35	-0.02	0.07	0.24	0.51	0.88	91,757
Measured CAPM α (%)	0.30	0.77	-0.46	-0.09	0.24	0.62	1.08	90,375

(b) Monthly sample

Notes: This table provides summary statistics for our main sample of funds. Panel (a) is based on the quarterly sample of funds that have a non-missing beta gap, as defined in Section 4.2. Inside ownership comes from form ADV. Leverage is defined as the ratio of gross-to-net assets. Share restrictions are a capital-weighted average of the number of days over which investors can redeem shares in each fund. Panel (b) summarizes monthly returns for the same set of funds. Measured CAPM β and α are based on rolling 36-month regressions of fund returns on the CRSP VW index, both in excess of the one-month riskless rate. See Section 2.2 for more details on data and variable construction.
Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.

Table 2: Response Rates and Reported Equity-Market Betas

	Equity	Rates	Currency	Credit	Vol	Def-Corp	Commodity	Def-ABS
Fraction R-T	0.47	0.39	0.41	0.35	0.28	0.15	0.25	0.13
Fraction R-NT	0.28	0.30	0.24	0.27	0.23	0.29	0.15	0.24
Fraction NR	0.25	0.31	0.36	0.37	0.49	0.56	0.60	0.64
Unique Fund Count	4,574							
Total Fund-Quarter Obs	74,523							

(a) Response Rates by Risk Factor Exposures

	Mean	p5	p25	p50	p75	p95	N
$\beta(-20)$	0.21	-0.05	0.00	0.00	0.35	1.00	56,793
$\beta(-5)$	0.22	0.00	0.00	0.00	0.40	1.00	56,926
$\beta(5)$	0.24	0.00	0.00	0.00	0.40	1.00	56,742
$\beta(20)$	0.24	0.00	0.00	0.00	0.40	1.00	56,859
$\beta(-20)/\beta(-5)$	0.96	0.50	1.00	1.00	1.00	1.25	51,778
$\beta(20)/\beta(5)$	0.99	0.75	1.00	1.00	1.00	1.25	51,310
$\beta(-20)/\beta(20)$	0.90	0.00	1.00	1.00	1.00	1.00	56,009

(b) Summary Statistics of Reported Equity-Market Betas

Notes: Table 2a shows the fraction of fund-quarter observations in which a fund reports the extent to which it is exposed to a particular risk factor. R-T means the fund reports the factor is both relevant and tested; R-NT means the fund reports it is relevant, but is not tested; and NR means the fund reports that factor is not a relevant risk factor in a given quarter. Def-ABS and Def-Corp respectively correspond changes in default rates for asset-backed securities and corporate bonds. Vol corresponds to changes in option-implied volatility. See Section 3.1.1 for details. Table 2b reports the distribution of equity factor betas implied by the reported equity factor exposures. The equity factor question gives four scenarios for returns to a diversified equity index: down 20%, down 5%, up 5%, and up 20%. Fund i reports the change, as a % of net assets, in response to each of these scenarios. From each scenario s, we can then construct an implied $\beta(s)$ as $\beta(s) = \%\Delta NAV/\%\Delta Market(s)$. Funds report the effect of each scenario separately for their long and short positions, and we net these effects first before calculating betas. Data for both panels is at the fund-quarter level.

Data Sources: U.S. Securities and Exchange Commission Form PF.

Table 3: Baseline Flow-Performance Regression

	$Flow_{i,q+1}$				
	(1)	(2)	(3)	(4)	(5)
$\widehat{\alpha}_{i,q}$	1.80***	1.29***	1.92***	1.86***	1.44***
	(7.46)	(6.69)	(7.22)	(7.48)	(6.84)
$(\beta_{i,q} - \widehat{\beta}_{i,q}) \times \mu^m_{i,q}$	-0.66**	-0.60***	-0.75**	-0.65*	-0.69**
	(-2.57)	(-2.68)	(-2.64)	(-1.96)	(-2.29)
ω	0.37	0.46	0.39	0.35	0.48
$t(\omega)$	2.61	2.56	2.76	1.91	2.18
Controls	No	Yes	No	No	Yes
Liquidity FE	No	No	Yes	No	Yes
Strategy-by-Time FE	No	No	No	Yes	Yes
Adjusted R^2	0.02	0.09	0.02	0.02	0.09
N	28,905	28,572	28,902	28,905	28,569

Notes: This table reports regressions of the following form:

$$flow_{i,q+1} = \theta_0 + \theta_1 \hat{\alpha}_{i,q} + \theta_2 \left(\beta^R_{i,q} - \widehat{\beta}_{i,q}\right) \times \mu^M_{i,q} + \Gamma X_{i,q} + \varepsilon_{i,q+1},$$

where $flow_{i,q+1}$ is the flow into fund i in quarter $q+1$, $\hat{\alpha}_{i,q}$ is the fund's realized CAPM alpha, measured from a 36-month rolling regression using returns through the end of quarter q. $\beta^R_{i,q}$ is the reported equity beta by fund i from the last month of quarter q and $\hat{\beta}_{i,q}$ is the estimated beta based on a 36-month rolling return regression. $\mu^M_{i,q}$ is the market return measured over the same 36-month window as $\hat{\alpha}_{i,q}$ and $\hat{\beta}_{i,q}$. $X_{i,q}$ is a vector of controls measured in quarter q that includes log net asset value and fund flows. Liquidity fixed effects are based on deciles for the weighted average number of days needed for investors to redeem capital. t-statistics are reported below regression estimates and are based on standard errors from Driscoll and Kraay (1998) with eight lags. The unit of observation in the regression is fund and quarter. See Section 4.2 for more details.

Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.

Table 4: Flow-Performance Regressions for Positive and Negative Beta Gap

	$Flow_{i,q+1}$							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\widehat{\alpha}_{i,q}$	1.57***	1.13***	2.00***	1.60***	2.07***	1.46***	1.62***	1.36***
	(6.28)	(5.60)	(5.68)	(4.80)	(7.44)	(5.55)	(6.51)	(6.18)
$(\beta_{i,q} - \widehat{\beta}_{i,q}) \times \mu^m_{i,q}$	0.13	0.10	-1.38***	-1.21***	0.31	0.38	-0.85***	-0.87***
	(0.26)	(0.22)	(-4.10)	(-5.06)	(0.52)	(0.54)	(-4.41)	(-3.50)
ω	-0.08	-0.09	0.69	0.76	-0.15	-0.26	0.53	0.64
$t(\omega)$	-0.26	-0.22	2.95	3.42	-0.50	-0.51	4.59	3.21
Subsample	β-Gap(+)	β-Gap(+)	β-Gap(−)	β-Gap(−)	Non-Linear β	Non-Linear β	Linear β	Linear β
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Liquidity FE	No	Yes	No	Yes	No	Yes	No	Yes
Strategy-by-Time FE	No	Yes	No	Yes	No	Yes	No	Yes
Adjusted R^2	0.01	0.08	0.02	0.09	0.03	0.09	0.01	0.08
N	11,458	11,329	17,447	17,240	5,300	5,250	19,532	19,272

Notes: This table reports subsample regressions of the same form contained in Table 3. In columns (1)-(4), subsamples are determined based on whether a fund has a positive or negative beta gap, $g_{i,q} = \beta^R_{i,q} - \hat{\beta}_{i,q}$, as of the end of quarter q. In columns (5)-(8), subsamples are determined based on whether the fund reports a non-linear beta as of the end of quarter q. t-statistics are reported below regression estimates and are based on standard errors from Driscoll and Kraay (1998) with eight lags. The unit of observation in the regression is fund and quarter. See Table 3, Section 4.2, and Section 4.3.3 for details.
Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.

Table 5: Flow-Performance Regressions for Inside Ownership

	$Flow_{i,q+1}$							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\widehat{\alpha}_{i,q}$	2.32***	1.87***	2.90***	2.34***	1.73***	1.29***	2.00***	1.52***
	(4.91)	(5.06)	(4.56)	(4.28)	(6.60)	(6.15)	(6.40)	(6.73)
$(\beta_{i,q} - \widehat{\beta}_{i,q}) \times \mu^m_{i,q}$	-0.34	-0.23	-0.62	-0.43	-1.17***	-1.04***	-1.09***	-1.19**
	(-1.06)	(-0.80)	(-0.76)	(-0.64)	(-3.21)	(-3.23)	(-2.98)	(-2.66)
ω	0.15	0.12	0.21	0.18	0.67	0.81	0.54	0.78
$t(\omega)$	1.18	0.85	0.86	0.71	2.69	2.54	2.81	2.52
Subsample	Low I-O	Low I-O	Low I-O	Low I-O	High I-O	High I-O	High I-O	High I-O
Excl. Insiders	No	No	Yes	Yes	No	No	Yes	Yes
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Liquidity FE	No	Yes	No	Yes	No	Yes	No	Yes
Strategy-by-Time FE	No	Yes	No	Yes	No	Yes	No	Yes
Adjusted R^2	0.02	0.10	0.01	0.05	0.02	0.08	0.01	0.05
N	13,871	13,698	13,598	13,238	13,967	13,823	13,888	13,687

(a) Inside Ownership

	$Flow_{i,q+1}$			
	(1)	(2)	(3)	(4)
$\widehat{\alpha}_{i,q}$	1.51***	1.99***	0.87***	1.57***
	(3.57)	(3.99)	(3.45)	(5.70)
$(\beta_{i,q} - \widehat{\beta}_{i,q}) \times \mu^m_{i,q}$	0.98*	-0.45	-0.11	-1.68***
	(1.83)	(-1.39)	(-0.27)	(-4.69)
ω	-0.65	0.22	0.12	1.07
$t(\omega)$	-1.46	1.33	0.28	3.40
I-O Subsample	Low I-O	Low I-O	High I-O	High I-O
Beta-Gap	+	-	+	-
Controls	Yes	Yes	Yes	Yes
Liquidity FE	Yes	Yes	Yes	Yes
Strategy-by-Time FE	Yes	Yes	Yes	Yes
Adjusted R^2	0.08	0.10	0.07	0.08
N	5,586	8,107	5,260	8,540

(b) Inside Ownership Interacted with the Beta Gap

Notes: This table reports subsample regressions of the same form contained in Table 3. Subsamples in Table 5a are determined based on whether a fund has above- or below-median inside ownership, as reported in Form ADV. In columns (3),(4), (7), and (8), we instead construct fund-level flows that exclude inside ownership. Subsamples in Table 5b are based on inside ownership and the sign of the beta gap, $g_{i,q} = \beta^R_{i,q} - \hat{\beta}_{i,q}$, in quarter q. t-statistics are reported below regression estimates and are based on standard errors from Driscoll and Kraay (1998) with eight lags. The unit of observation in all regressions is fund and quarter. See Table 3, Section 4.2, and Section 4.3.4 for complete details.

Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.

Table 6: Fees and Strategic Communication

	Fee_i			
	(1)	(2)	(3)	(4)
$\widehat{\alpha}_i$	0.24***	0.35***	0.17***	0.32***
	(14.62)	(9.26)	(5.80)	(9.55)
$(\bar{\beta}_i^R - \widehat{\beta}_i) \times \mu_i^m$	-0.06	-0.32***	0.03	-0.32**
	(-0.62)	(-2.90)	(0.26)	(-2.54)
Subsample	β-Gap(+)	β-Gap(-)	Non-Linear β	Linear β
Liquidity FE	Yes	Yes	Yes	Yes
Strategy FE	Yes	Yes	Yes	Yes
Size FE	Yes	Yes	Yes	Yes
Adjusted R^2	0.79	0.87	0.70	0.84
N	827	1,206	289	1,744

Notes: This table shows subsample regressions of the average realized fee of each fund on its unconditional realized alpha and its beta gap times the average realized market risk premium. Average realized fees are computed as the difference between each fund's gross- and net-of-fee returns. Realized alpha is based on a full sample regression of gross returns on the CRSP VW index, both in excess of the riskless rate. Columns (1) and (2) are subsample regressions for funds with a positive and negative beta gap, respectively. Columns (3) and (4) are subsample regressions for funds with a non-linear or linear perceived beta, respectively. Whether a fund has a non-linear beta is defined unconditionally based on whether fund i's manager reports a non-linear beta $O_{i,q}$ for at least half of its sample. See Section 4.4 and Section 4.3.3 for details. t-statistics are reported below regression estimates and are based on standard errors clustered by adviser.

Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.

Table 7: Are investors surprised by fund performance?

	Abnormal Flow$_{i,2020Q1}$				
	(1)	(2)	(3)	(4)	(5)
$1\{\beta_i - \widehat{\beta}_i > 0\}$	-0.43	-0.45	-0.49	-1.02	-1.21
	(-0.43)	(-0.45)	(-0.49)	(-1.01)	(-0.90)
$(\beta_i - \widehat{\beta}_i)$	6.83**	6.95**	7.03**	5.64**	2.22
	(2.24)	(2.28)	(2.29)	(2.04)	(0.88)
$1\{\beta_i - \widehat{\beta}_i > 0\} \times (\beta_i - \widehat{\beta}_i)$	-14.92***	-15.09***	-15.18***	-12.76***	-7.60*
	(-3.56)	(-3.59)	(-3.61)	(-2.98)	(-1.70)
Approach	Baseline	(s, l, h, g)	m	m	Fund
FE	None	None	None	(s, l)	None
R^2	0.04	0.04	0.04	0.18	0.01
N	889	889	889	887	825

Notes: This table shows regressions of the following form:

$$\varepsilon_{i,2020Q1} = a + b_1 1(g_i > 0) + b_2 g_i + b_3 1(g_i > 0) g_i + \nu_i,$$

where $\varepsilon_{i,2020Q1}$ is the abnormal flow of fund i in 2020Q1, g_i is its beta gap estimated using data up to 2020, and $1(g_i > 0)$ is an indicator for whether the beta gap is positive. Abnormal flows $\varepsilon_{i,2020Q1}$ are constructed using a two-stage procedure. In stage 1, we use data prior to 2020 to estimate a regression of flows on lagged performance. In stage 2, we take the estimated coefficients from stage 1 and use them to compute the expected flow in 2020Q1. Abnormal flows are then the difference between realized flows in 2020Q1 and the predicted flow form Stage 2. Columns (1)-(5) differ in the exact procedure used to construct abnormal flows in stage 1. Column (1) assumes the same performance-flow relationship for all funds. Column (2) lets the relationship vary by fund strategy, liquidity, size, and the sign of the beta gap. Columns (3) and (4) allow it to vary by manager. The difference between the two is that column (4) includes a strategy-by-liquidity fixed effect in the regression of abnormal flows on the beta gap. Column (5) lets the performance-flow relationship vary by fund. See Section 5.1 for details.

Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.

Table 8: Predictors of the β Gap

	Adj. R^2 (%)	p(FE=0)	p(FE equal)	N
Fund FE	54	0.00	0.00	2,552
Adviser FE	40	0.00	0.00	3,495
Strategy FE	2	0.00	0.00	3,649
Year FE	1	0.00	0.00	3,649

(a) Variance Decomposition of β Gap

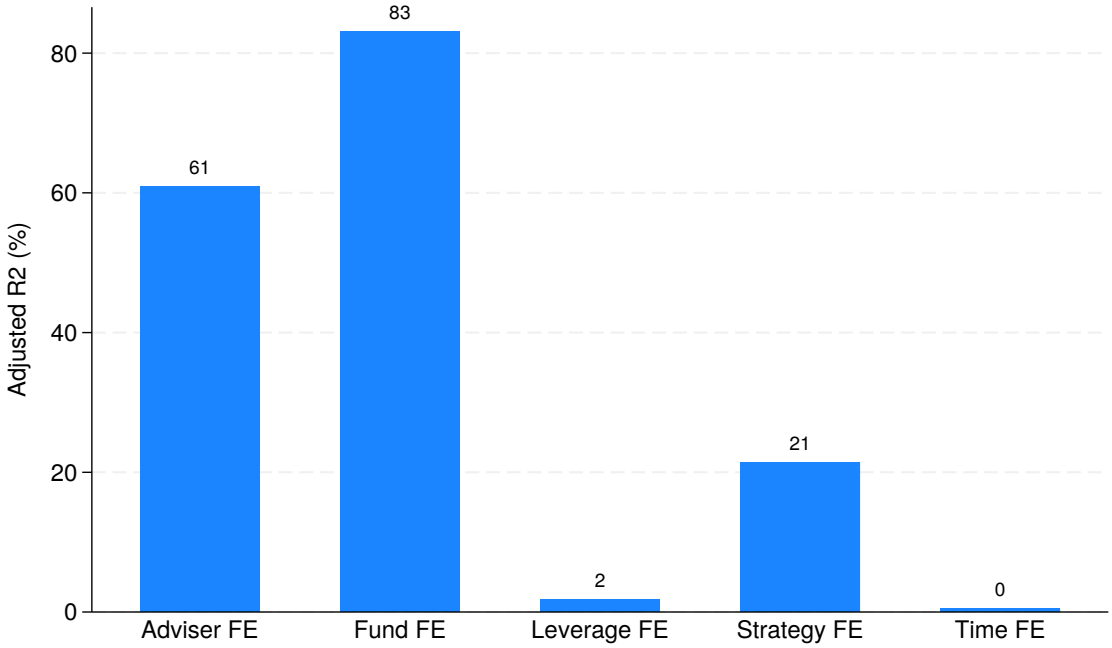
	$100 \times 1\{\beta_{i,t} - \widehat{\beta}_{i,t} > 0\}$				
	(1)	(2)	(3)	(4)	(5)
$\widehat{\beta}_{i,t-1}$	-9.98***				
	(-8.51)				
1{High Leverage}		6.24***			
		(3.38)			
Derivatives Usage			2.09**		
			(2.06)		
1{Large Fund}				5.61***	
				(2.92)	
1{High-Return Fund}					8.92***
					(4.82)
R^2	0.04	0.00	0.00	0.00	0.01
N	1,972	3,573	2,179	3,649	3,649

(b) Predictive regressions

Notes: Panel (a) of the table shows regressions of the beta gap on various fixed effects. The statistics reported are the adjusted R^2, the p-value from a test of zero fixed effects, the p-value from a test of equal fixed effects, and the sample size. The fixed effect drops all singletons. Panel (b) shows regressions of an indicator for whether a fund has a positive beta gap on the beta implied by returns, an indicator variable for whether the fund is in the top tercile of gross-to-net assets (leverage), the ratio of gross notional in derivatives to NAV, an indicator variable for whether the fund is in the top tercile of size, and an indicator variable for whether the fund is in the top tercile of trailing 36-month returns. The dependent variable in this regression is multiplied by 100. Standard errors in panel (b) are clustered by fund and all continuous variables are standardized to have mean zero. The sample for all subtables is the same as in Figure 3 and ensures that each fund has non-overlapping samples.

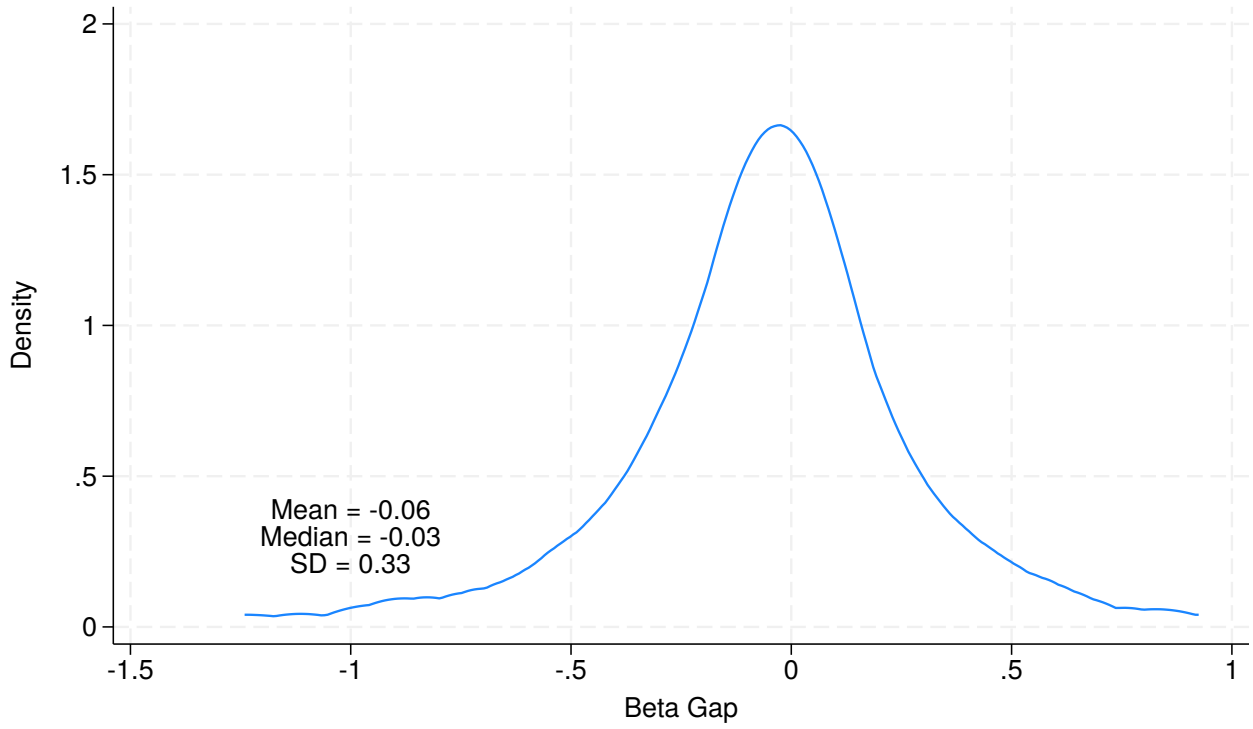
Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.

Figure 1: Variance Decomposition of Reported β



Notes: This figure shows the adjusted R^2 from a regression of each fund's perceived beta on various fixed effects. At the end of each quarter, each fund manager reports on Form PF its exposure to movements in equity markets. In each month, we combine the realized return of the CRSP VW index and the manager's last available response to construct the fund's perceived beta. See Section 3.1.1 for details.
Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.

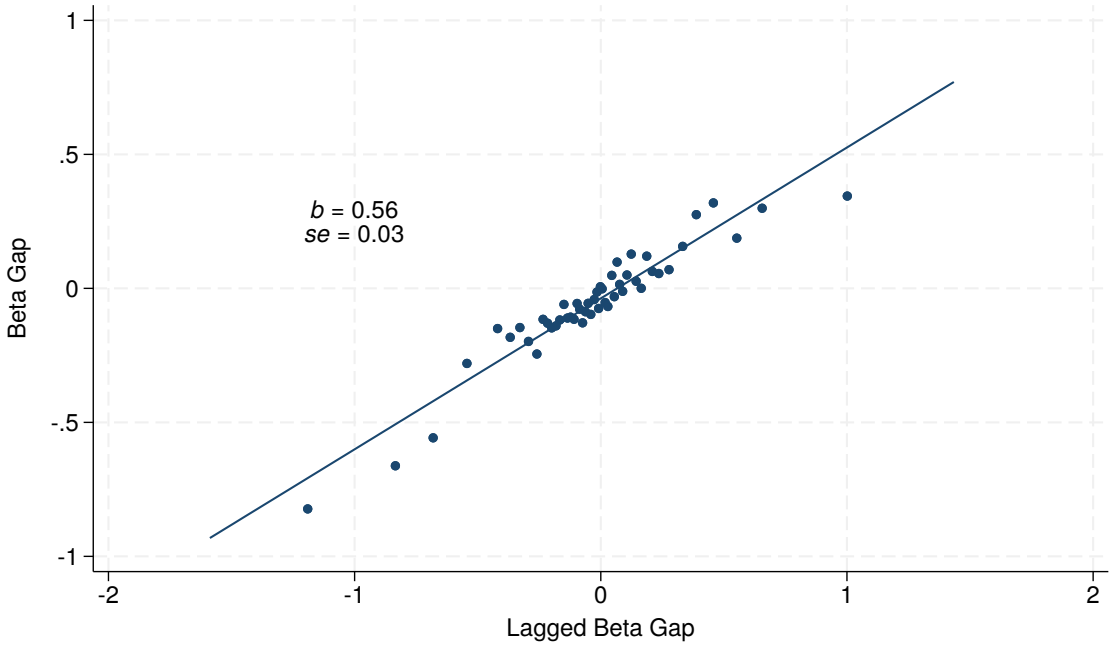
Figure 2: The Distribution of the Beta Gap



Notes: This figure shows a kernel density plot of the unconditional beta gap, $\bar{\beta}_i^R - \widehat{\beta}_i$. $\bar{\beta}_i^R$ is defined as the average monthly perceived beta over each fund's full sample. $\widehat{\beta}_i$ is estimated using a full-sample regression of each fund's monthly returns on the CRSP VW index, both in excess of the riskless rate. See Section 3.2.1 for details.

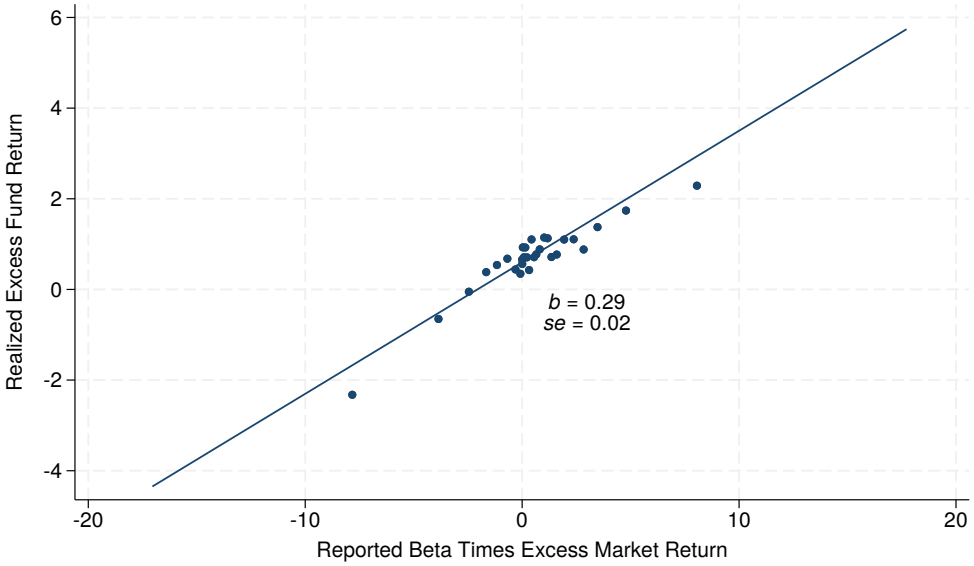
Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.

Figure 3: Is the Beta Gap Persistent?

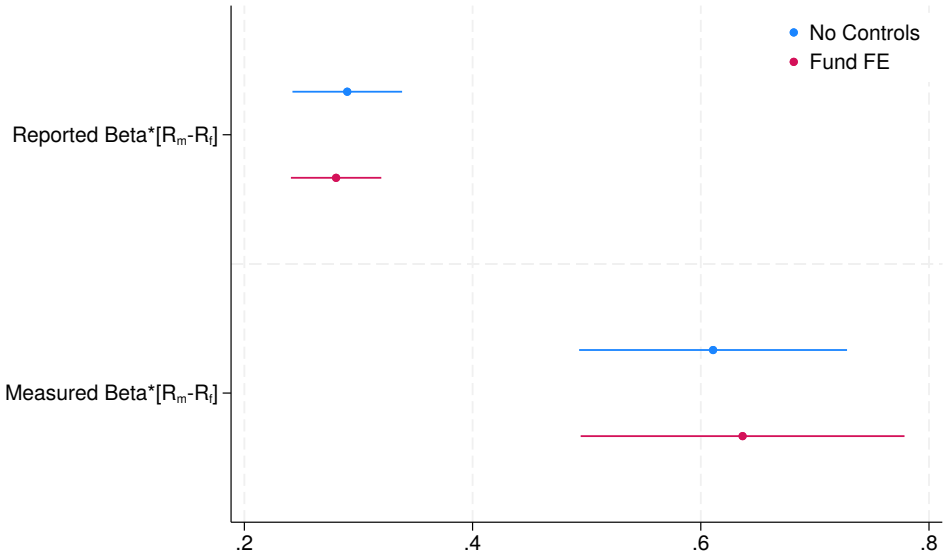


Notes: This plot shows a binned scatter plot of each fund's beta gap on its lagged value, using non-overlapping windows. For each fund, we construct 36-month non-overlapping windows and compute the beta gap based on data in each window. The binscatter then shows each fund's beta gap in a given window on its beta gap from the previous window. An OLS regression line is also included in the plot with standard errors based on clustering at the fund level.
Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.

Figure 4: Are Manager Perceptions Informative?



(a) Predictive Power of Manager-Perceived Beta

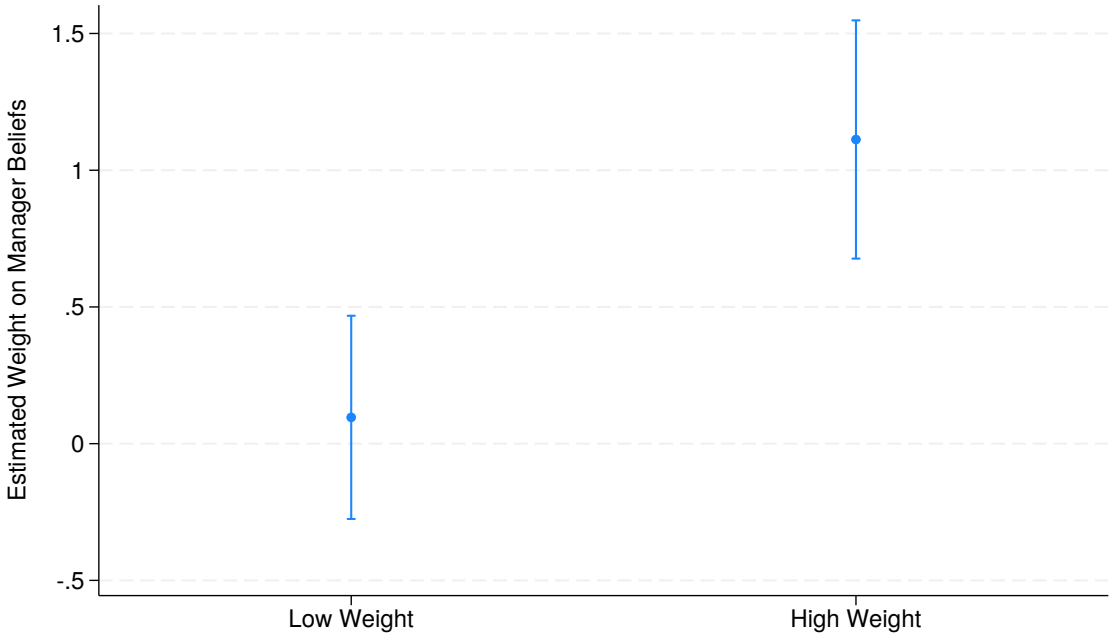


(b) Full Regression Results

Notes: Panel (a) of this figure shows a binned scatter plot of each fund's excess return $\tilde{R}_{i,t+1}$ in month $t+1$ on its manager's expected beta for the same month times the realized excess return of the CRSP-VW index, $\beta^R_{i,t+1}\tilde{R}_{M,t+1}$. The binned scatter plot controls for the measured beta using a 36-month rolling window through month t times the realized excess market return in month $t+1$, denoted by $\widehat{\beta}_{i,t}\tilde{R}_{M,t+1}$. The slope in the plot is based on an OLS regression with standard errors clustered by fund and month. Panel (b) shows estimated coefficients of OLS regressions of $\tilde{R}_{i,t+1}$ on $\beta^R_{i,t+1}\tilde{R}_{M,t+1}$ and $\widehat{\beta}_{i,t}\tilde{R}_{M,t+1}$. The blue lines are coefficients when no fixed effect is included and the red lines are coefficients when including a fund fixed effect. All standard errors and their implied confidence bands based on Driscoll and Kraay (1998) with fourteen lags.

Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.

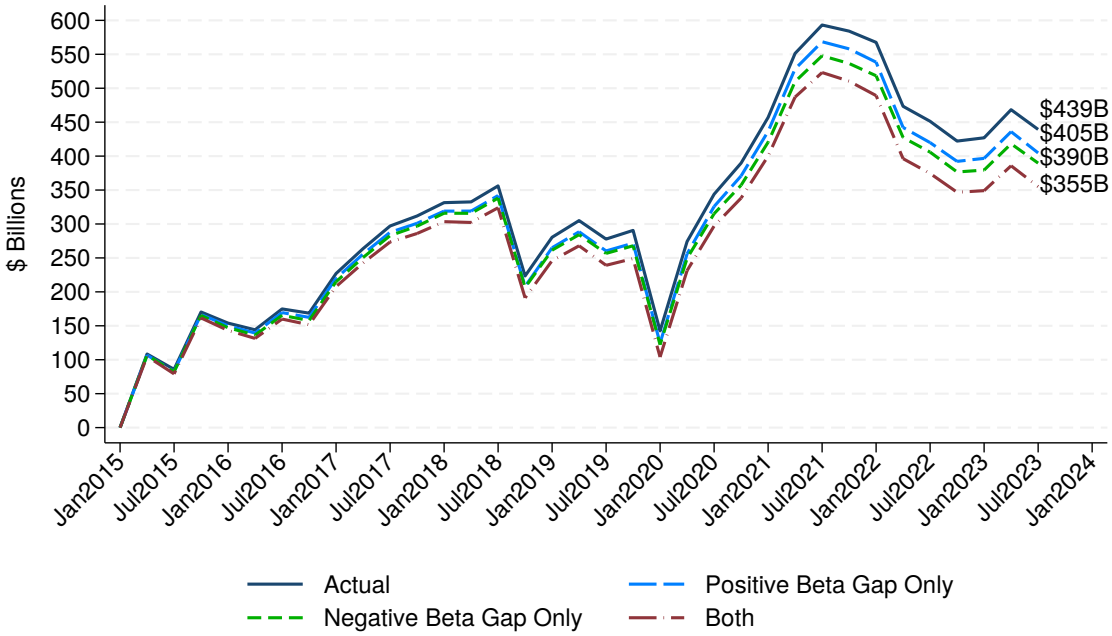
Figure 5: Heterogeneity in ω



Notes: This plot visualizes heterogeneity across funds in the implied weight that investors put on manager risk perceptions when risk-adjusting returns. To construct the plot, we estimate Equation (7) for each fund, constructing each fund's implied weight $\omega_f = -\theta_2/\theta_1$ from the estimated regression coefficients θ_1 and θ_2. We then sort funds based on whether they have above- or below-median ω_f. Within the set of the low-ω funds, we re-estimate Equation (7) and compute the average implied ω. We repeat this procedure for the set of high-ω funds and plot the resulting ω's, along with confidence bands that are constructed using the delta method. See Section 4.2 for complete details.

Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.

Figure 6: AUM under Counterfactual Flows



Notes: This figure shows actual aggregate flows into the sample of hedge funds for which we observe a beta gap from 2015 to 2023 (dark solid blue line). In addition, it shows three counterfactual sets of aggregate flows. In the first counterfactual (light blue dashed line), we construct flows into positive beta-gap funds under the assumption that investors put a weight of one on the manager's perceived risk when risk-adjusting returns. Flows into negative beta-gap funds are held constant at their true value. In the second counterfactual, we construct flows into negative beta-gap funds under the assumption that, when risk-adjusting returns, investors put the zero weight on the manager's perceived risk. Flows into positive-beta gap funds are held at their true value. In the last counterfactual, we assume investors in positive beta gap funds put full weight on manager risk perceptions and investors in negative beta-gap funds put zero weight on manager risk perceptions. See Section 6.1.

Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.

A Robustness Checks

A.1 Are Manager Perceptions of Positive Beta-Gap Funds Informative?

For the subsample of funds with a positive beta gap, Figure A1 shows that manager perceptions of risk are informative about future returns, even after controlling for the risk reflected in past returns.

A.2 Flow-Performance in Positive and Negative Beta-Gap Funds

Table A1 shows that our main conclusions regarding positive vs negative beta-gap funds are robust to:

1. Including fund and time fixed effects
2. Clustering by fund and time, as opposed to using Driscoll and Kraay (1998) standard errors as we do in the main text.
3. Using 36-month rolling averages of manager risk perceptions to define the beta gap (and its sign). This is the same window over which CAPM betas are estimated from returns.
4. Including funds that report they have zero betas with respect to all other factors. This alleviates concerns that some of the beta gap arises because managers report multivariate CAPM betas (i.e., holding other factors fixed) whereas we estimate univariate CAPM betas.

A.3 Learning During the Pandemic

In the main text, we estimate abnormal flows during the onset of Covid-19 using a two-step procedure. Here, we instead using a single panel regression to test if outflows were more sensitive to the beta-gap during the pandemic. Specifically, we estimate:

$$\begin{aligned} flow_{i,q} = \lambda_{s,l,h} &+ \sum_{j=1}^{4} \theta_j flow_{i,q-j} + \sum_{j=1}^{4} \kappa_j R_{i,q-j} \\ &+ \zeta_1 1(g_i > 0) + \zeta_2 g_i + \zeta_3 g_i 1(g_i > 0) \\ &+ C_q \left[\rho_0 + \rho_1 1(g_i > 0) + \rho_2 g_i + \rho_3 g_i 1(g_i > 0)\right] + \varepsilon_{i,q}, \end{aligned}$$

where $flow_{i,q}$ and $R_{i,q}$ are, respectively, the flow and return of fund i in quarter q. g_i is the beta gap for fund i, estimated using data through 2019Q4. $1(g_i > 0)$ is an indicator for whether the beta gap is positive. C_q is an indicator for whether quarter q equals 2020Q1. Table A2 shows the estimated coefficients for ρ_1, ρ_2, and ρ_3. They are comparable and tell the same story as our two-step procedure in the main text.

A.4 Predictors of the Beta Gap

In the main text, we predict the sign of the beta gap using fund characteristics that are presumably observable to investors. For completeness, Table A3 repeats the analysis using the beta gap itself, as opposed to its sign. The main conclusions are largely unchanged.

Table A1: Robustness: Flow-Performance Regressions for Positive and Negative Beta Gap

	$Flow_{i,q+1}$							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\widehat{\alpha}_{i,q}$	2.58***	1.13***	1.64***	1.19***	2.86***	1.60***	1.21***	1.74***
	(5.75)	(4.17)	(6.07)	(3.97)	(7.71)	(6.85)	(3.85)	(6.05)
$(\beta_{i,q} - \widehat{\beta}_{i,q}) \times \mu^m_{i,q}$	-0.08	0.10	3.14**		-1.33***	-1.21***	-1.39***	
	(-0.12)	(0.18)	(2.36)		(-2.77)	(-4.14)	(-4.98)	
$(\bar{\beta}_{i,q} - \widehat{\beta}_{i,q}) \times \mu^m_{i,q}$				-0.71				-1.68***
				(-1.02)				(-3.77)
ω	0.03	-0.09	-1.91	0.59	0.47	0.76	1.15	0.97
$t(\omega)$	0.12	-0.18	-1.88	1.00	2.54	3.83	3.30	3.83
Subsample	β-Gap(+)	β-Gap(+)	β-Gap(+)	β-Gap(+)	β-Gap(−)	β-Gap(−)	β-Gap(−)	β-Gap(−)
ZBOF only	No	No	Yes	No	No	No	Yes	No
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fund-and-Time FE	Yes	No	No	No	Yes	No	No	No
Liquidity FE	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Strategy-by-Time FE	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Adjusted R^2	0.02	0.10	0.08	0.11	0.03	0.11	0.15	0.10
N	11,191	11,329	2,142	9,949	17,059	17,240	4,429	13,655

Notes: This table reports subsample regressions of the same form contained in Table 3. Subsamples are determined based on whether a fund has a positive or negative beta gap, $g_{i,q} = \beta^R_{i,q} - \hat{\beta}_{i,q}$, as of the end of quarter q. Columns (1)-(4) focus on funds with a positive beta gap. In column (1), we include fixed effects for fund and time. In column (2), we include fixed effects for strategy-by-time and liquidity. In column (3), we restrict the sample to funds that report zero beta on all non-equity factors (ZBOF). In column (4), we define the beta gap (and its sign), using rolling 36-month windows of each fund manager's perceived beta, $\beta^R_{i,q}$. Columns (5)-(8) repeat the specifications but for funds with a negative beta gap. t-statistics are reported below regression estimates. In columns (2) and (6), standard errors are clustered by fund and quarter. For all other columns, standard errors are based on Driscoll and Kraay (1998) with eight lags. The unit of observation in the regression is fund and quarter. See Table 3, Section 4.2, and Section 4.3.3 for details.

Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.

Table A2: Robustness: Are investors surprised by fund performance?

	$\text{Flow}_{i,q}$			
	(1)	(2)	(3)	(4)
$1\{\beta_i - \widehat{\beta}_i > 0\}$	-0.95***	-0.46	-0.42	-0.87**
	(-4.32)	(-1.44)	(-1.23)	(-2.73)
$(\beta_i - \widehat{\beta}_i)$	6.76***	6.10***	5.53***	5.60***
	(11.65)	(10.78)	(7.45)	(6.38)
$1\{\beta_i - \widehat{\beta}_i > 0\} \times (\beta_i - \widehat{\beta}_i)$	-12.84***	-13.53***	-12.74***	-13.63***
	(-12.84)	(-11.76)	(-11.08)	(-12.86)
Approach	Baseline	(s, l, h)	(s, l, h, g)	m
R^2	0.12	0.10	0.10	0.08
N	19,412	19,412	19,369	19,422

Notes: This table shows regressions of the following form:

$$\begin{aligned} flow_{i,q} = \lambda_{s,l,h} + \sum_{j=1}^{4} \theta_j flow_{i,q-j} + \sum_{j=1}^{4} \kappa_j R_{i,q-j} \\ + \zeta_1 1(g_i > 0) + \zeta_2 g_i + \zeta_3 g_i 1(g_i > 0) \\ + C_q \left[\rho_0 + \rho_1 1(g_i > 0) + \rho_2 g_i + \rho_3 g_i 1(g_i > 0)\right] + \varepsilon_{i,q}, \end{aligned}$$

where $flow_{i,q}$ and $R_{i,q}$ are, respectively, the flow and return of fund i in quarter q. g_i is the beta gap for fund i, estimated using data through 2019Q4. $1(g_i > 0)$ is an indicator for whether the beta gap is positive. C_q is an indicator for whether quarter q equals 2020Q1. The table shows the estimated coefficients for ρ_1, ρ_2 and ρ_3. In column (1), the fixed effect $\lambda_{s,l,h}$ is based on the intersection of fund strategy, declines of share liquidity, and deciles of size. Column (2) retains this fixed effect, includes only one lag of returns, then allows the slope on lagged returns to vary by strategy-liquidity-size. Column (3) allows both the fixed effect and coefficient on lagged return to vary with strategy, liquidity, size, and the sign of the beta gap. Column (4) allows both to vary by adviser. The regression is estimated using all data through 2020Q1. t-statistics are listed below point estimates are based on standard errors are from Driscoll and Kraay (1998) with eight lags. See Sections 5.1 and A.3 for details.

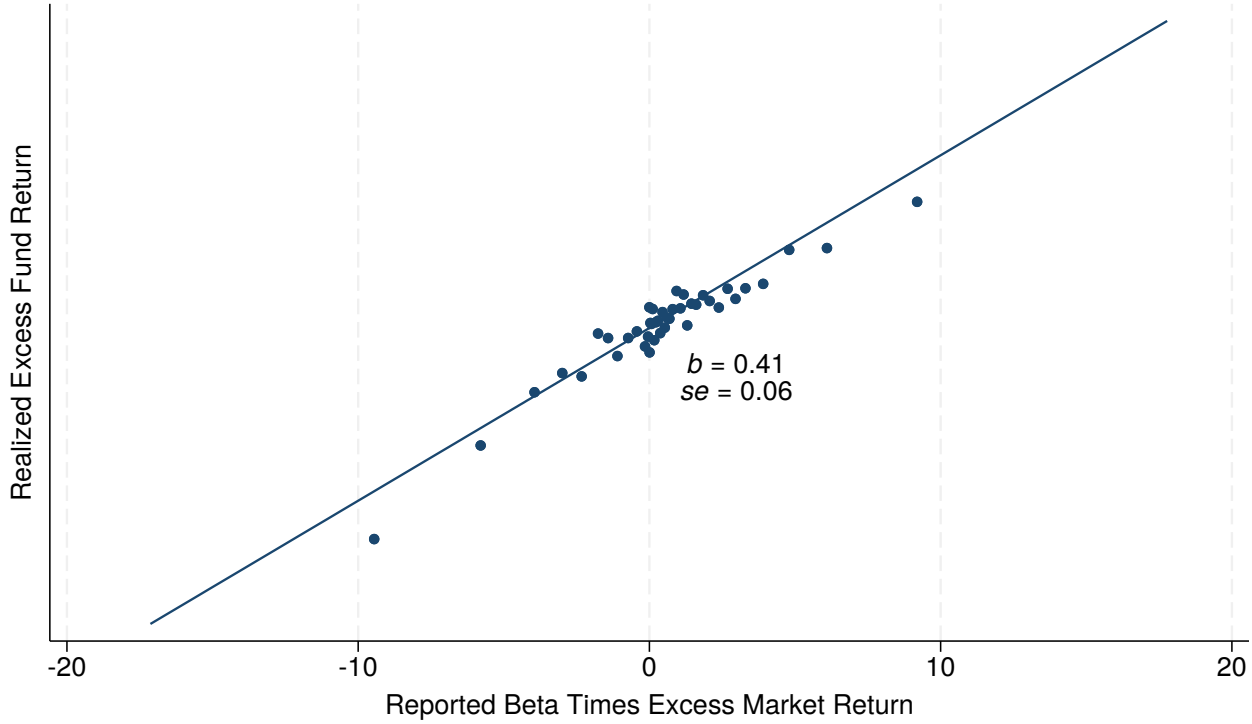
Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.

Table A3: Robustness: Predictors of the β Gap

	$\beta_{i,t} - \widehat{\beta}_{i,t}$				
	(1)	(2)	(3)	(4)	(5)
$\widehat{\beta}_{i,t-1}$	-0.18***				
	(-14.79)				
1{High Leverage}		0.09***			
		(6.40)			
Derivatives Usage			0.03***		
			(4.41)		
1{Large Fund}				0.01	
				(0.48)	
1{High-Return Fund}					0.01
					(0.42)
R^2	0.22	0.01	0.01	0.00	0.00
N	1,972	3,573	2,179	3,649	3,649

Notes: This table shows regressions of the beta gap on the beta implied by returns, an indicator variable for whether the fund is in the top tercile of gross-to-net assets (leverage), the ratio of gross notional in derivatives to NAV, an indicator variable for whether the fund is in the top tercile of size, and an indicator variable for whether the fund is in the top tercile of trailing 36-month returns. Standard errors are clustered by fund and all continuous regressors are standardized to have mean zero. The sample for this table is the same as in Figure 3 and ensures that each fund has non-overlapping samples.
Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.

Figure A1: Are Manager Perceptions of Positive Beta-Gap Funds Informative?



Notes: This figure shows a binned scatter plot of each fund's excess return $\tilde{R}_{i,t+1}$ in month $t+1$ on its manager's expected beta for the same month times the realized excess return of the CRSP-VW index, $\beta^R_{i,t+1}\tilde{R}_{M,t+1}$. The binned scatter plot controls for the measured beta using a 36-month rolling window through month t times the realized excess market return in month $t+1$, denoted by $\widehat{\beta}_{i,t}\tilde{R}_{M,t+1}$. The slope in the plot are based on an OLS regression with standard errors clustered by fund and month. Only funds with a positive beta gap are included.

Data Sources: U.S. Securities and Exchange Commission Form PF; Center for Research in Security Prices, CRSP 1925 US Indices Database, Wharton Research Data Services, http://www.whartonwrds.com/datasets/crsp/.