

Limits of Arbitrage under the Microscope: Evidence from detailed Hedge Fund Transaction Data

Bastian von Beschwitz*
Federal Reserve Board

Sandro Lunghi**
Inalytics

Daniel Schmidt***
HEC Paris

March 15, 2017

Abstract

We exploit detailed transaction and position data for a sample of long-short equity hedge funds to document new facts about the trading activity of sophisticated investors. We find that the initiation of both long and short positions is associated with significant abnormal returns, suggesting that the hedge funds in our sample possess investment skill. In contrast, the closing of long and short positions is followed by return continuation, implying that hedge funds close their positions too early and “leave money on the table.” As we demonstrate with a simple model, this behaviour can be explained by hedge funds being (risk) capital constrained and facing position monitoring costs. Consistent with our model, we document that the return continuation following closing orders is more pronounced when these constraints become more binding (e.g., after negative fund returns or increases in volatility).

JEL classification: G11, G12, G14, G15

Keywords: Hedge funds, Short selling, Profitability, Limits of arbitrage

* Bastian von Beschwitz, Federal Reserve Board, International Finance Division, 20th Street and Constitution Avenue N.W., Washington, D.C. 20551, USA, tel. +1 202 475 6330, e-mail: bastian.vonbeschwitz@gmail.com (corresponding author).

** Sandro Lunghi, Inalytics, 9th Floor, Corinthian House, 17 Lansdowne Road, Croydon CR0 2BX, UK, tel. +44 (0)20 3675 2904, e-mail: alunghi@inalytics.com.

*** Daniel Schmidt, HEC Paris, 1 Rue de la Libération, 78350 Jouy-en-Josas, France, tel. +33 (0)139 67 9408, e-mail: schmidt@hec.fr.

We thank Chris Collins and Laura Kane for excellent research assistance and Inalytics Ltd. for providing the data. We thank Adam Reed and seminar participants at McGill, the Federal Reserve Board, and the 9th Annual Hedge Fund and Private Equity Research Conference for helpful comments. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

Fundamental arbitrage, defined as the process of acquiring an information advantage through fundamental research and trading on it, plays a critical role for market efficiency as it helps to align individual stock prices with their (ever elusive) “fair” values.¹ In this paper, we analyze proprietary transaction data for a sample of discretionary long-short equity hedge funds to shed light on this important activity. Indeed, these type of hedge funds routinely undertake independent long and short investments (“directional bets”), making them the archetype of a fundamental arbitrageur. Exploiting the richness of our data, we uncover new stylized facts pertaining to the *initiation and closure of long and short* arbitrage positions, and discuss how these findings relate to the broader research on informed trading, short selling, and the limits of arbitrage. We view our paper as providing first-hand *micro*-evidence for the importance of several channels highlighted in these literatures.

Our data comprises the *entire trading history as well as daily position updates* for 21 long-short hedge funds over a ten-year period. This level of detail allows us to conduct a microscopic analysis of hedge funds’ trading activity. In particular, we are able to distinguish buy transactions that initiate a long position (“long buys”) from buys that close an existing short position (“short buys”). Similarly, we distinguish sells that initiate a short position (“short sells”) from sells that close an existing long position (“long sells”). We begin our investigation with an examination of the profitability of these different trades. Focusing on the initiation and closure of trading positions, we find that long buys and short sells—i.e., trades that open new long and short positions—are, respectively, followed by significantly positive and negative benchmark-adjusted returns with an absolute magnitude of about 1.5% (2%) over the next 125 (250) trading days. This proves that the hedge funds in our sample possess investment skill.

In stark contrast, we find that closing trades are not informed. To the contrary, long sells and short buys tend to be followed, respectively, by positive and negative returns; that is, returns in the opposite direction of the closing trade. When we design a trading strategy that goes long in stocks in which hedge funds just closed a long position (long sells) and shorts stocks from closed short positions (short buys) we obtain a significant alpha of about 1% over the next six months (125 trading days). This implies that the hedge funds in our sample close their positions too early in the sense that they “leave money on the table.”

To understand this finding, we draw on a simple trading model (presented in the appendix) in which a hedge fund decides whether and how much to invest in mispriced stocks. We embed three important—and as we

¹ Fundamental arbitrage is different from relative-value arbitrage; i.e., the process of exploiting price differences between assets or portfolios that provide (almost) identical payoffs in the future. As it violates “the law of one price,” arbitrage opportunities of this kind provide a striking failure of market efficiency, perhaps explaining why they have received more attention in the academic literature. Yet, eliminating short-term price discrepancies is not the same as having more informative prices (see Brunnermeier (2005), and Weller (2016)) and fundamental arbitrage should be especially important toward the latter goal.

believe realistic—features into the model: First, the hedge fund is assumed to face a risk constraint, which prevents it from taking too large a position in any mispriced stock. Second, the hedge fund incurs a monitoring cost for each open position in its portfolio. The first assumption mirrors standard practice in the hedge fund industry (see, for instance, Pedersen (2015)) and should thus be uncontroversial. The second assumption can be interpreted loosely as a fixed transaction, monitoring or attention cost for maintaining the position and checking whether a previous trading signal has not lost its allure. Such a cost naturally leads the investor to focus on a limited number of open positions, consistent with what we empirically observe for the long-short equity hedge funds in our sample.² Finally, we assume that new investment opportunities (stock mispricings) emerge each period and partially decay over time. This is again consistent with what we find in the data as we show that roughly 75% (=1.5%/2%) of the alpha is earned within the first six months after the initiation of a stock position. We describe the hedge fund’s optimal trading rule (and characterize it formally in the appendix).

We show that, under these assumptions, the hedge fund’s optimal trading rule involves early position closures in order to free up risk capital and/or monitoring capacity for new trading opportunities. This insight explains our main finding that hedge funds close positions that would continue to generate alpha going forward. Furthermore, the model makes additional predictions: First, at any point in time, the profits from newly opened positions should exceed the profits that hedge funds forego by closing existing ones. Second, the return continuation following position closures should be more pronounced—meaning that the hedge fund leaves more money on the table—when the fund (1) simultaneously opens a lot of new positions (that require capital), (2) has suffered from poor past performance, and (3) when the risk constraint becomes more binding due to a surge in fund return volatility.

We test and confirm these predictions in the data. We begin by comparing returns after initiating and closing orders within the same month. We find that, over the 60 trading days following the order, newly initiated long (short) positions yield an alpha that is 0.6% larger (smaller) than the alpha following closed long (short) positions. Thus, we document that, within the same month, hedge funds generate more alpha with their opening trades than they forego by closing their trades prematurely. This shows that hedge funds indeed recycle their limited risk capital into more profitable trading opportunities.

Next, we conduct a number of sample splits for the trading strategy built around hedge funds’ closing trades—i.e., going long (short) in stocks from closed long (short) positions. This strategy can be thought of as measuring how much alpha hedge funds forego by closing early. First, we examine whether this strategy is more profitable when hedge funds have higher opportunity costs due to facing more trading opportunities. To proxy for the change in their trading opportunities, we look at whether hedge funds increase or decrease

² For instance, the average hedge fund in our sample has less than 80 open positions at any time.

the number of open portfolio positions. We find that the strategy yields a highly significant excess return of 2.5% over the next half year after an increase in the number of positions, while we observe only a return of 0.6% after a decrease in the number of open positions. Second, we conduct a similar sample split based on whether the fund had a positive or a negative return over the prior week. The idea of this test is that negative returns reduce the available (risk) capital of the fund, forcing it to close down some existing stock positions. Indeed, we find that trading against closing orders following a negative fund return yields an excess return of 2.8% over the next half year, while trading against closing orders following a positive fund return only delivers 0.4%. Third, we split the sample based on whether the fund experienced an increase in volatility. Since we posit that hedge funds operate under a risk constraint, we expect such an increase in volatility to result in additional position closures. Consistent with this intuition, we find that trading against closing orders following an increase in fund volatility yields an excess return of 2.0% over the next half year, while it gives only 0.9% after a decrease in fund volatility.

Taken together, these findings suggest that early position closures—which are associated with foregone risk-adjusted returns—occur in response to binding (risk) capital and attention constraints, rather than being due to hedge funds' ignorance and/or investment biases. To the best of our knowledge, the results thus offer some of the first micro-level evidence for several themes commonly discussed in the limits of arbitrage literature—in particular, recent strands that emphasize the importance of (risk) capital, leverage constraints and limited attention (see, e.g., Shleifer and Vishny (1997), Gromb and Vayanos (2010)).

As mentioned before, one of the key advantages of our data is that we are able to cleanly identify the complete set of long and short positions of long-short equity hedge funds. This allows us to study the performance differences between the closing and opening of positions and also relate them to the past returns of the fund. Moreover, we argue that long-short equity hedge funds constitute an ideal laboratory to examine the limits to *fundamental* arbitrage: First, long-short equity hedge funds take *directional* bets on individual stocks rather than engaging in relative value arbitrage. Thus, the long and short positions are not part of a combined trading strategy as in, for example, merger arbitrage, allowing us to directly examine the profits of individual positions. Second, discretionary long-short equity hedge funds spend a lot of effort and money on research to distill an information advantage from public and private data sources (Pedersen (2015)). In other words, they closely resemble the textbook case of informed arbitrageurs trying to trade on fundamental mispricing.

While the availability of complete trading records and portfolio holdings are clear advantages of our data, we acknowledge that the relatively small number of hedge funds (21) raises some questions about the selection and representativeness of our sample. We try to allay such concerns to the best of our ability. First, we note that the funds in our sample represent a variety of different sizes and, like many long-short equity funds, invest in equity markets worldwide. On many dimensions, we have no reasons to believe that the funds in our

sample are markedly different from a typical long-short equity hedge fund. Second, as we explain in the robustness section, the nature of our data make it unlikely that common sample problems—such as survivorship bias or back-filling bias—should be of major concern for us. Finally, we emphasize that all the results are rationalized with the help of our stylized trading model. Thus, to the degree that our model—and in particular the assumptions on risk limits and position monitoring costs—capture features that are common in the hedge fund industry, one would expect our results to be generalizable to the broader population.

In conclusion, our results make for a rich description of the anatomy of hedge fund trading. They show that long-short equity hedge funds in our sample are skilled but constrained investors: While their opening trades are clearly profitable (both on the long and short side), they do not hold on to their positions until the alpha is fully exploited. Rather, they close their positions prematurely in order to recycle their capital and/or accommodate tightened risk constraints. Thus, the hedge funds in our data behave like constrained arbitrageurs as they are portrayed in the limits of arbitrage literature. By trading on their positive and negative information, they reduce mispricings and contribute to market efficiency. At the same time, their constraints expose them to external pressures (stemming from, e.g., portfolio losses or risk management considerations) that can force them to close their arbitrage positions. While we do not document this here, it is clear that this behavior can have grave consequences for price stability and market efficiency. For instance, when many funds are forced to close their positions at the same time, their very behavior can render mispricings more severe and leave market prices destabilized. More generally, our results show how mispricings can persist in the market despite informed arbitrageurs.

Our paper contributes to several strands of research. First and foremost, we speak to the literature on hedge funds—and in particular on their trading and performance. Many papers examine hedge fund skill using databases on self-reported hedge fund returns, but are hampered by different biases of these databases (see, e.g., Agarwal, Mullally, and Naik (2015) for a survey). More recent papers try to examine hedge fund skill using quarterly 13F filings data and reach mixed conclusions: While Cao et al. (2016) find that hedge fund holdings predict future returns, Griffin and Xu (2009) find no predictive power of hedge fund holdings. Jank and Smajlbegovic (2015) use data from hedge funds' mandatory disclosure of large short positions and find evidence for predictability. We add to the debate on hedge fund performance by examining trading skill using *complete* trading and position records for *both* long and short positions of the same funds. Furthermore, our data has the advantage of covering *all* equity positions of the hedge funds in the sample. We find strong evidence of hedge fund outperformance for up to one year after the opening of positions. This shows that long-short equity funds in our sample possess the skill to identify mispriced stocks, thereby complementing previous work that emphasize hedge funds' role as liquidity providers (Jylhä, Rinne and Suominen (2014), Franzoni and Plazzi (2015), Jame (2016)). Finally, our work is closely related to Choi, Pearson and Sandy (2016), who study hedge fund short positions gleaned from merging institutional transaction data from

ANcerno with quarterly holdings from 13F. They find that the position openings by hedge funds in their sample do *not* predict long-term returns. Instead, they find that their short positions are profitable only over the short-term (up to 5 trading days), suggesting that these funds make the bulk of their profits from liquidity provision. Our data is arguably more comprehensive³ and, more importantly, seems to cover the trading activity for a *different class* of hedge funds—long-short equity—as our results for long-term predictability are markedly different.

Second, we contribute to the literature on short selling. Several papers find that short selling predicts future returns (e.g. Desai, Thiagarajan, and Balachandran (2002), Boehmer, Jones, and Zhang (2008), Diether, Lee, and Werner (2009), Asquith, Pathak, and Ritter (2005), Engelberg, Reed, and Ringgenberg (2012)). However these papers usually focus only on short selling or the change in short interest. We add to these papers by examining the profitability of both the *initiation* and *closing* of short positions. We find negative returns after short sells, but no positive returns after short buys—suggesting that only the opening trades for new short positions are informed. The only other paper examining returns following the closing of short positions is Boehmer, Duong, and Huszar (2015). Their results differ from ours in that they show evidence of positive return predictability for closing trades. However, their analysis is based on the mandatory disclosure of very large position closures and may thus be influenced by price impact and signaling effects.

Third, we contribute to the literature on informed trading and the limits of arbitrage. Using transaction data for a sample of active mutual funds from the same data provider, Di Mascio, Lines and Naik (2016) provide an in-depth analysis on the opening of long arbitrage positions. Like us, they find abnormal returns of about 1.5% following the opening of a long position (but over a longer horizon of 18 months). They further document that the mutual funds gradually build up their positions over a horizon of up to 6 months, consistent with them trying to strategically limit their price impact. While we find less of this behaviour in our sample,⁴ we complement their study by including short sales and by focusing on and explaining the returns following position closures. We thereby link with the (mostly theoretical) literature on the limits of arbitrage that emphasizes different channels as to why arbitrageurs may be forced to close their positions too early. Our findings offer empirical support for many of these channels, and thus provide important micro-level evidence for their validity.

The remainder of this paper is organized as follows. Section I describes the simple trading model we have in mind and lay out its testable predictions. Section II describes the data and provides summary statistics. Section III focuses on the profitability of the initiation and closing of long and short positions. In Section IV, we relate

³ We have access to daily as opposed to quarterly position updates and ANcerno only seems to cover a subset of the stock trades undertaken by hedge funds contained in that sample (see, e.g., Di Mascio, Lines, and Naik (2016)).

⁴ On average, opening orders in our sample already make up about 77% of the maximum position size that we observe for the position.

post-closure returns to several proxies of hedge funds' shadow cost of capital. Section V provides robustness checks. Section VI concludes with a discussion of the broader implications of our findings for market efficiency and the limits of arbitrage.

I. Hypotheses

We think of the long-short equity hedge funds in our data as fundamental arbitrageurs. Rather than holding a diversified portfolio to earn the risk premium, they collect and analyze public and private data to form an assessment about the fundamental value of a specific target company and establish a long (short) position when they find the company's stock to be sufficiently under- (over-)valued. The starting point of our empirical investigation is therefore to see whether the long and short stock positions opened by hedge funds in our sample deliver risk-adjusted returns (alpha). Prior research on hedge fund performance and managerial skill are hampered by data constraints and reach mixed conclusions (see, for instance, the survey by Agarwal, Mullally, and Naik (2015)). Given that our data, while covering only 21 funds, is the most detailed hedge fund transaction and position data so far studied in the academic literature, our performance analysis constitutes a valuable contribution in its own right.

Next, we investigate how and when hedge funds close their positions. If hedge funds were unconstrained, we would expect them to hold on to their positions until all the alpha is reaped, implying that post-closure risk-adjusted returns should be zero. In practice, however, we expect hedge funds to be capital and/or attention constrained: while they can take leverage, their ability to do so depends on banks willingness to provide it, and the fundamental research required to identify and monitor mispriced stocks should be time and effort intensive, implying that long-short equity hedge funds focus on a limited number of open positions (directional bets). Being cognizant of their constraints, we expect hedge funds to allocate their limited resources on the basis of a cost-benefit analysis. Each period, they will decide how many positions to maintain, which ones to open and which ones to close. An important implication is that, when constraints are binding, the hedge fund may decide to close an arbitrage position before its alpha is fully exploited. Thus, if our sample hedge funds are indeed capital or attention constrained as we posit, we expect the returns of long (short) positions to remain positive (negative) even after these positions have been closed. These post-closure returns should, however, be smaller than the returns following newly opened positions—for otherwise the hedge fund would have been better off holding on to the old position.

To guide our intuition as to when early position closures should occur, we develop and solve a simple trading model in which a hedge fund faces a risk constraint, incurs position monitoring costs, and new stock mispricings pop up every period but gradually decay over time. In Appendix B, we describe our model in

detail and derive the hedge fund's optimal trading rule. Here, we summarize its key intuitions and the resulting empirical predictions.

Our modeling assumptions are supposed to reflect realistic features of the trading environment for long-short equity hedge funds: The risk constraint is meant to capture, in a simplified way, common risk management practices such as risk parity investment (see Pedersen (2015)). A straightforward implication of this constraint is that position sizes are bounded and inversely related to the volatility of the underlying stock.⁵ The position monitoring cost is a placeholder for any type of fixed cost that is associated with holding a stock position. For instance, it can represent a fixed transaction cost or a fixed attention cost for monitoring a given position (the hedge fund may want to check, for example, whether the trading signal, which induced the opening of the position, is still valid after the arrival of new information). Without this assumption, the hedge fund would always smoothly scale back position sizes all the way to zero until the alpha is fully exploited. Thus, there would be no early position closures. With a fixed monitoring cost, early position closures do occur as it is not economical to hold on to a position below a certain minimum position size. A natural implication of this assumption is that wealthier funds have more open positions—a prediction for which we find strong support in the data.⁶

Our model identifies three potential reasons for why a hedge fund may close a position before its alpha is fully exploited: First, because the fund only maintains a limited number of open positions, it may close some positions when better investment opportunities arise. Second, as the hedge fund's wealth decreases, the monitoring of existing positions becomes more expensive relative to downscaled position sizes, to which it responds by reducing the total number of open positions. Third, as the hedge fund's stock positions become more volatile, it is forced to downscale its positions in order to satisfy the risk limit constraint. This again increases the monitoring costs *per dollar invested* and thus leads to a reduction in the number of total positions. Hence, the model predicts that the closure of long (short) positions should be followed by more positive (negative) returns when the hedge fund (1) simultaneously opens new positions (as a proxy for having many new investment opportunities), (2) has had a poor past performance, and (3) when the funds' stock positions become more volatile.

In the next section, we introduce our data before proceeding to our empirical tests of these predictions.

⁵ The precise nature of the risk constraint is not important. For example, we obtain similar predictions if we replace it by a leverage constraint as modeled in Gromb and Vayanos (2002). The only important thing is that, since both long and short positions consume from the overall risk budget (or margin capital), the hedge fund is prevented from leveraging up his positions without limit.

⁶ When we split hedge funds into above and below median in terms of total portfolio value, we find that small and large funds have, respectively, 61 and 94 open stock positions, on average.

II. Data and Variable Construction

A. *Inalytics data*

Our data on long-short equity hedge funds is provided by Inalytics Ltd. and this is the first time it is used in an academic collaboration. A different subset of the Inalytics database, for long only equity funds, has been previously studied in Di Mascio, Lines and Naik (2016). Inalytics analyses patterns of investment decisions to interrogate the robustness of various investment processes. The output is portfolio monitoring for institutional asset owners and investment consultants and process improvement for asset and hedge fund managers. When an institutional client investor, like a plan sponsor, makes an allocation to a hedge fund, this hedge fund provides its' trading and portfolio data to Inalytics so that they can monitor the fund's performance on behalf of their client.

Our data contain *complete* trading and portfolio information for the equity holdings of 21 hedge funds. For each fund, we are thus able to track both their long and short portfolios. Specifically, we have access to two datasets: The first is a trade-level dataset containing all trades. Variables in this dataset include stock identifiers (ISIN, SEDOL, and CUSIP), the date of the trade, the number of shares traded, and the execution price. The second dataset is a stock-day level dataset of each funds' portfolio holdings. This dataset contains stock identifiers, the number of shares held, and the price of the stock at the end of the day. All prices are expressed in the base currency of the fund and in the local currency of the stock.

We use a merged dataset that combines the holdings and trading data (details on merging these two datasets can be found in Appendix A.3). Hedge funds often split their orders into several trades that are executed on different days to reduce the market impact of their orders. To avoid double counting, we follow Di Mascio, Lines, and Naik (2016) and aggregate the trades belonging to one investment decision into orders. We assume that trades belong to one order if they trade the same stock in the same direction and the distance between two trades of the order is two days or less. Seventy-three percent of the orders consist of only one trade.

B. *Summary statistics*

Our sample period runs from 2005 to 2015. However, each individual fund covers only a fraction of this sample period. Figure 1 gives an overview about how the number of funds, positions and orders changes over the sample period. From 2005 to 2007, the sample is fairly small with only 1-6 funds. From late 2008 to mid-2013, we have 8 to 9 funds in the sample. In 2013, the number of funds jumps to 17. However, the early funds have more positions, so from 2008 Q1 to the end of the sample period we always have at least 500 positions in the data. Orders move more proportional to the number of funds. From 2008 Q1 to the end we have around 20 orders per day in the sample, but towards the end of the sample period that number jumps to over 100

orders per day. We include our full sample period in our tests to preserve statistical power and ensure that no specific time period is driving our results.

In Table 1, Panel A, we display summary statistics by fund. Funds hold on average 50 long positions and 24 short positions (median values are 36 and 19). The fewer number of short positions is also reflected by the fact that short positions make up about 30% by USD value. Having a larger long than short portfolio is seen as typical for long-short equity hedge funds (Pedersen (2015)). The funds conduct on average 6 orders per day. Compared to an average of 74 positions this corresponds to a new order for a stock every 15 trading days. The daily USD trading volume of a fund divided by its holdings is on average 5.4% (median 2.8%). Our funds have fairly different sizes. The 10th percentile fund holds 0.12 billion USD in assets, while the 90th percentile fund holds 6 billion USD. The median fund holds 0.4 billion USD. Also the investment areas of the funds differ. We present them in Figure 2. We have 7 Europe-focused funds, 3 US, 3 UK and 2 Australia-focused funds, as well as 6 funds that invest world-wide. Due to that, European stocks are somewhat overweighted in our sample of fund holdings (they make up 29% of stocks held, compared to making up 17% of world market capitalization in Worldscope).

In Panel B, we display summary statistics by position. A position lasts from its initiation—i.e., the first buy for long positions or the first sell for short positions—to its close—i.e., the moment when the stock holding goes back to zero. After being closed, a new position can be established in the same stock. However, this does not happen very often: on average there are only 2 positions in a given stock over the lifetime of the fund. Our data contains about 16,000 positions; 6.9% of them are already open when the fund enters the database, while 11% are still open when the fund leaves the database. Due to this censoring, the length of positions will be biased downwards. The investment horizon of the funds seems to be fairly long. The average length of a position is 104 trading days (about half a year), but the median is only 35 trading days (about 2 months). Over the lifetime of a position, funds conduct on average 6 orders (median 3) and change the direction of trading on average 2.5 times (median 1).

Next, we examine summary statistics at the order-level. We distinguish between three types of orders: Initial orders that initiate the position, closing orders that close the position and follow-up orders that change the size of the position in between. We display summary statistics for each type of order separately in Panels C to E. The initial and closing orders are much larger than the follow-up orders: when standardized by the maximum size of a given position, initial and closing orders on average make up around 77% of this maximum position size (median 100), while the follow-up orders make up only 15.8% (median 8.7%). This suggests that the important investment decisions are the initiations and closures of positions. Follow-up orders are small and thus more likely to be based more on liquidity or rebalancing motives rather than on information. Still, the follow-up orders make up 69% of all orders in our sample. Hedge funds don't split

orders into separate trades very often: the average number of trades per order is only about 1.6 and the median is 1 for each order type.

C. Datastream and Worldscope data

Because the hedge funds in our sample trade stocks internationally, we require international stock market and balance sheet data. We use the datasets most commonly used in the international context: Datastream for stock returns and Worldscope for balance sheet data. For stocks that are in the Inalytics data but are not covered in Datastream we add stock return information provided by Inalytics (this affects 14% of our stocks). We use this data mainly to compute risk-adjusted returns to benchmark the performance of different funds and orders. We use three methodologies to risk-adjust returns: we compute excess returns with respect to the fund-specified benchmark, we compute risk-adjusted returns following the methodology of Daniel, Grinblatt, Titman, and Wermers (1997), hereafter DGTW, and we compute alphas using the four-factor model of Carhart (1997). The details of the risk-adjustments are explained in Appendix A.3; here we provide only a brief summary description.

Excess returns are computed as returns minus the return of the fund-specified benchmark. This means that the risk-adjustment depends on the fund, i.e. the excess return of the same stock may be positive for one fund but negative for another fund. The benchmarks can even be different within the same investment area. For example, some Europe-focused funds benchmark against the MSCI Europe, while others benchmark against the FTSE Europe. However, the benchmarks are always the same for both long and short positions of the same fund and they do not change over time.

As a second methodology, we compute DGTW returns on a regional level. We split the world into 5 regions (Japan, North America, Europe, Asia-Pacific and Emerging Markets) following Karolyi and Wu (2014). The assignment of countries into regions is displayed in Appendix A.2. Within each region, we sort stocks into quintiles by market capitalization, market-to-book ratio and past-12 months returns, thus forming 625 portfolios (125 per region). We compute DGTW returns as stock returns minus the returns of the respective portfolio. We choose to compute portfolios on a regional level because it is more granular than using the same portfolios for the whole world and prior evidence suggests that local factors are better in pricing risk (Griffin (2002)). At the same time, many countries do not have enough stocks to populate 125 portfolios. Working at the regional level therefore appears to be a reasonable compromise.

As a third methodology, we implement a regional version of the Carhart (1997) 4-factor model. For each of our 5 regions we compute a market factor, a High minus Low Book to Market Factor (HML), a Small minus Big (SMB) factor and a Momentum (MOM) factor of winners minus losers. For each stock, we compute betas with respect to these factors using daily regressions over the prior 12 months. We then shrink these betas to

their cross-sectional average following Vasicek (1973). This implementation follows the suggestions of Levi and Welch (2016), who find that for predicting betas, daily regressions are more accurate than monthly regressions, 1 year horizons are better than longer horizons and it is important to shrink the betas. Finally, we compute alphas on the daily level as:

$$\text{Four factor } \alpha_{c,t} = r_{c,t} - \beta_c * (r_{m,t} - r_{f,t}) - \beta_{HML} * HML_t - \beta_{SMB} * SMB_t - \beta_{MOM} * MOM_t$$

More details on the construction of the risk-adjustment is presented in Appendix A.3.

III. Profitability Results

A. Profitability of opening and closing trades

We display gross fund profitability computed from holdings by year in Figure 3. In Panel A, we display the actual profitability of the fund. Because most funds have more long than short positions, this profitability co-moves a lot with the market. The worst year is 2008 when equity markets crashed worldwide in the wake of the Lehman bankruptcy. In 2009, equity markets recovered and our sample hedge funds see their best year. To get a better idea of the fund’s stock-picking skill, we display profitability based on equal-weighting the long and short portfolios in Panel B. Now 2008 is actually the best year and the worst year is 2009. This year is generally known as a bad year for hedge funds because of the so called momentum crash (Daniel and Moskowitz (2013)). With the exception of 2009, the funds always exhibit positive returns that are fairly stable in the 2-8% range. This seems to suggest that the funds in our sample exhibit skill.

We now examine this point in more detail by studying the post-trade returns for stocks bought and sold by hedge funds in our sample. We start with a simple graphical analysis presented in Figure 4. We show cumulative returns in excess of the fund-specified benchmark in the 250 trading days (approximately 1 year) following an order. We include only orders that either initiate or close a position (that is, we exclude follow-up orders). We split orders by whether they are related to a long or a short position.

Figure 4 reveals clear evidence of informed trading in the initiation of positions: in the 250 days following the initiation of a long (short) position, cumulative excess returns are approaching 2% (-2%). Moreover, on both the long and the short side, a large fraction of these returns is realized in the first 125 trading days (6 months) following the initiating order, after which the return drift appears to be a lot more muted. In other words, the post-trade alphas (per unit of time) for initiated positions decay over time: they are highest immediately after the position is established and then gradually shrink as time progresses.⁷

⁷ This finding is consistent with the evidence provided in Di Mascio, Lines and Naik (2016) who document a similar convexity in the cumulative abnormal returns following the initiation of long positions by the long-only mutual funds in their sample.

In contrast, the closing of long and short positions does not seem to be informed. Long sells are not followed by negative returns, but rather by positive returns. In the 250 days following the closing of a long position cumulative excess returns are slightly below 1%. Similarly, short buys are followed by negative excess returns (-1% after 250 days). In both cases, most of the cumulative return is realized in the first 125 trading days following the order.

Next, we investigate the statistical significance of these findings. In Table 2, we focus on position initiations and run a regression of risk-adjusted returns following the order on $D(Short)$, a dummy variable equal to one if the order initiates a short position (and zero if it initiates a long position). We examine all three measures of risk-adjusted returns for holding periods of 60 and 125 trading days (approximately 3 and 6 months) following the order. We measure returns from the date following the last date of the order; i.e., we do not account for within order returns. We include fund fixed effects to control for any differences in post-trade profitability across funds that could correlate with their propensity to enter a short position. We also include month fixed effects to ensure that our results are not driven by a particular time period. Finally, we cluster standard errors two-way by stock and last date of order. Clustering by stock accounts for correlation due to overlapping returns (for example if two orders happen in the same stock in the same month) and clustering by date accounts for correlation in the cross-section of stock returns.

Given our specification, the coefficient estimate for the $D(Short)$ dummy can be interpreted as the return difference between long and short positions that have been initiated in the same month. The results, presented in Table 2, show that this return difference is economically and statistically significant. For instance, for excess returns, long positions outperform short positions by about 1.9% over 60 days and 2.6% over 125 days. For DGTW returns and alphas the effect is slightly smaller at about 1.6% over 60 days and 2% over 125 days. These results are all statistically significant at the 1% level. In short, hedge funds' position openings are clearly profitable, suggesting that the fund managers in our sample possess investment skill.

In Table 3, we repeat our analysis for closing orders. We again find a negative coefficient for the $D(Short)$ dummy, albeit with a smaller economic magnitude. For excess returns, the return difference between closed long and closed short positions equals 0.7% over 60 days and 1.5% over 125 days. For DGTW returns and alphas the effect is again slightly smaller. Over the 125 days horizon, the return difference is statistically significant for all measures of risk-adjusted returns. These results suggest that the hedge funds in our sample close their positions too early in the sense that these positions would have earned significant risk-adjusted returns going forward. However, such early position closures don't have to be suboptimal. Indeed, we argue below that they can be explained by the presence of risk capital and/or attention constraints faced by the hedge funds in our sample.

These results have important implications for our understanding of the informativeness of different types of trades. Indeed, at least for the long-short equity hedge funds in our sample, only the trades that initiate a new stock position are informative, whereas those that close a position are not only uninformative but rather predict returns in the opposite direction of the closing trade. Our analysis thus shows that, in order to study the informativeness of individual buy and sell trades (such as attempted in Franzoni and Plazzi (2015) or Jame (2016)), it becomes important to determine whether these trades open or close a stock position, which is only possible with access to portfolio data such as we use here. Without this distinction, long buys and short buys as well as short sells and long sells are lumped together, causing a significant downward bias when assessing the profitability of these trades.

B. Opening a new stock position vs. holding-on to an old one

We have established that both the opening and the closure of a long (short) position is followed by positive (negative) returns. As argued in the hypotheses section, a natural explanation for this is the presence of a risk capital (or margin capital) constraint: a constrained hedge fund may want to close an existing stock position even though it still offers some alpha in order to free-up capital that can be invested into new, more promising, trading opportunities. Of course, this argument only makes sense when these new investments indeed deliver higher returns than those that are foregone by closing existing positions. A casual inspection of Figure 4 suggests that this is indeed the case: newly established positions earn most of their alpha in the first weeks/months after the opening trade. After some time, alphas peter out and so it could be more attractive to open a new position.

We now test for this prediction more rigorously. Specifically, we check in a regression setting whether post-initiation returns outperform post-closure returns after controlling for month and fund fixed effects. This approach allows us to compare initiations and closures that occur roughly at the same point in time—and where it is thus conceivable that the closure provided the capital for the new position opening. The key variable of interest is $D(Initiation)$, a dummy variable that takes the value one when the order initiates a (long or short) position and zero when it closes the position (follow-up orders are again excluded from this analysis).

Table 4 shows the results. In Panel A, we focus on long positions only. The significantly positive coefficient for the $D(Initiation)$ dummy implies that newly initiated long positions are indeed more profitable than the previous long positions that are closed within the same month by about 0.6-1% depending on the risk-adjustment and the holding horizon. For short positions (Panel B), the coefficient flips sign, meaning that initiated short positions are followed by more negative returns than closed short positions (although it is not always significant). In Panel C, we examine both long and short positions jointly. To be able to combine long and short positions, we use signed returns instead of returns as the dependent variable. Signed returns are defined as returns for long positions and minus one times the returns for short positions. We find about 0.6-

0.7% higher signed returns following the initiation of positions. Because combining short and long positions improves statistical power, these tests are all highly statistically significant. This finding shows that hedge funds are on average right when they reallocate their capital into new stock positions.

In summary, the results of this section show that the hedge funds in our sample possess investment skill but face constraints: they initiate stock positions that generate alpha, but close them before this alpha is fully exploited so that they are able to recycle their capital into new investment opportunities. In the next section, we investigate position closures in greater detail.

IV. Explaining Post-Closure Returns

In Appendix B, we show with the help of a stylized trading model that early position closures can be explained by funds being subject to risk capital constraints and position monitoring costs. In this section, we wish to provide further support for this mechanism by testing three distinct predictions from our model.

The first prediction says that existing stock positions should be closed earlier at times when more new trading opportunities emerge, and when the hedge fund thus opens many new stock positions. A larger number of early position closures in turn implies that hedge funds “leave more money on the table”—i.e., the return difference between closed long and closed short positions should grow. In Table 5, we test this prediction by splitting the sample of closing orders by whether the hedge fund increased or decreased the number of open positions over the previous 5 days (Panel A) or over the previous 10 days (Panel B). We then repeat our regression analysis from Table 3 for these different subsamples. Table 5 shows the results for a holding period of 125 trading days.⁸ The results broadly confirm our prediction: whereas the excess return difference between closed long and short positions after increases in the number of open positions over the previous 5 days is 2.5%, it is only 0.6% and insignificant after decreases in the number of open positions over the previous 5 days. The results for DGTW returns and 4-Factor Alpha are similar with 1.6% vs. 0.5%, and 2.0% vs. 0%. Furthermore, the results are robust to using the change over the previous 10 trading days instead of the previous 5 trading days (Table 5, Panel B). In summary, our results suggest that early position closures appear to be more common when hedge funds simultaneously seize new trading opportunities.

The second prediction concerns the relation between past portfolio profits and subsequent position closures. In the model, the hedge fund’s optimal number of open positions is pinned down, among other things, by the fund’s equity wealth (or total net asset value) relative to the position monitoring cost. Intuitively, this fixed cost makes it uneconomical to hold positions below a certain minimum position size. As such, wealthier funds

⁸ The results for 60 days go in the same direction but are economically weaker.

naturally hold a larger number of open positions, and when a given fund suffers portfolio losses it may respond by closing existing positions. We thus check whether the returns from the post-closure investment strategy from Table 3 are more pronounced after times in which the fund has experienced negative portfolio returns. The results, shown in Table 6, support this prediction. When we split closing orders by prior fund returns over the previous five trading days, the excess return difference between closed long and short positions is 2.8% in the subsample with negative prior fund returns and only 0.4% in the subsample of positive prior fund returns. The results are similar when we split the sample based on fund returns over the previous 10 trading days instead. These findings suggest that trading losses force funds to close some of their positions earlier, leaving more money on the table.

The third prediction follows from the risk constraint: when the volatility of stock returns goes up, hedge funds have to curb their position sizes in order to satisfy their risk constraint. Because of the fixed position monitoring cost, this may again cause the premature closure of some outstanding stock positions. We therefore conduct a sample split by the change in fund return volatility compared to the previous month.⁹ When volatility goes up, the volatility constraint becomes tighter and the fund may be forced to reduce the number of open positions, leading to earlier closure of positions. In contrast, when volatility decreases, there is less need for early position closures. Table 7 confirms this prediction. Focusing on excess returns over a 125-days horizon, we see that the return difference between closed long and short positions amounts to a statistically significant 2% at times when fund volatility goes up, while it is less than 1% and insignificant when volatility goes down. For DGTW returns (but not for 4-factor alphas), we find a similar difference between the two subsamples.

Overall, the findings in this section are consistent with a model in which hedge funds close their positions due to risk capital constraints and position monitoring costs. In other words, the hedge funds in our sample resemble constrained arbitrageurs as they are portrayed in the limits to arbitrage literature.

V. Selection Concerns and Robustness Checks

A) *Data issues and selection concerns*

Our data sample is quite small in that it covers only 21 long-short equity hedge funds. In our view, this drawback should be easily outweighed by the unprecedented level of detail that this data provides. Indeed, to the best of our knowledge, our data constitutes the most detailed and first *complete* account of the stock

⁹ We compute monthly fund return volatilities as the standard deviation of daily fund returns. We set volatility to missing when we have fewer than 15 non-missing daily return observations for a given month. We then calculate monthly fund return volatility changes.

trading activity for a sample of hedge funds studied in the academic literature to date.¹⁰ In this section, we discuss potential data biases and selection concerns that one could hold toward this data.

We begin by noting that several sample biases that have been identified in the literature should not be of major concern here. For instance, since hedge funds that engage with Inalytics provide their transaction data in real time, there should be little incentives for window dressing and back-filling bias should be limited. Moreover, since our data also covers funds that have already been terminated, survivorship bias should not be an issue.

One important remaining concern with our data is self-selection into the sample. It is not clear, however, which way the selection would go. On the one hand, successful hedge funds may be keen to engage with Inalytics in order to advertise their trading success. On the other hand, institutional clients may demand from a hedge fund to submit their trades to Inalytics for monitoring and verification purposes—especially when the hedge fund has performed poorly. To the extent that these biases roughly balance each other, one can thus think of our hedge funds as being more or less representative for their class (long-short equity).

Finally, we would like to point out that any remaining selection concerns should only affect the inference about the representativeness of the magnitude of the average trade performance that we document. They should not, however, cause us to question the validity of the limits to arbitrage channels for which we find evidence. Indeed, we believe that these channels should apply to every arbitrageur that faces some sort of risk-capital or margin-capital constraint. Given the general paucity of data on arbitrageurs' detailed activities, we are just the first to be able to provide direct evidence for them.

B) Examining all orders

In Section III, we analyze the profitability of initiating and closing orders. That is, follow-up orders are excluded from this analysis, even though they make up about 69% of all trade orders in our sample. This choice is motivated by the intuition that, of all orders, initiating orders should be the most informed (as they capture the point in time when a hedge fund started acting on its trading signal), whereas closing orders should in principle be the least informed (as an unconstrained hedge fund will only close after fully exploiting its trading signal). For follow-up orders, it is less clear what to expect. On the one hand, hedge funds may gradually build-up their arbitrage positions so as to minimize their price impact, in which case their follow-

¹⁰ The data that comes closest to ours in its level of detail is obtained via a fuzzy name-matching approach between the hedge fund trades contained in the ANcerno institutional transaction data and their quarterly equity holdings reported in 13F filings. However, funds covered by ANcerno only make available a subset of their transaction records and identifying long and short positions—a crucial element of our study—from quarterly holdings is bound to be noisy. Finally, while our data is on long-short equity hedge funds (for which we document clear evidence of stock-picking ability), the hedge funds covered in ANcerno appear to belong to a different class of hedge funds as they make most of their profits from liquidity provision and their trades do not predict long-term alpha (Franzoni and Plazzi (2015), Jame (2016), Choi, Pearson and Sandy (2016)).

up orders would appear to be informed (see Kyle (1985), Foster and Viswanathan (1996), Di Mascio, Lines and Naik (2016)). On the other hand, hedge funds may want to rebalance their positions to a specific target portfolio weight, implying that they gradually realize some of their trading profits over time, making follow-up orders look the opposite of informed.

Fortunately, these two cases are easy to distinguish: when a follow-up order increases an existing position (in absolute terms), it should tend to be informed. In contrast, when it decreases the position, it is likely to be a rebalancing trade which looks uninformed. Based on this intuition, we can repeat our analyses from Tables 2 and 3 after including follow-up trades. Specifically, in Table 8 Panel A, we focus on all orders that initiate or increase a given stock position (long buys and short sells) and test whether long positions continue to outperform short positions. We find that this is the case: returns after orders that increase a long position are larger by 1.2-2.0% over 60 days and 1.6-3.1% over 125 days, depending on the risk adjustment, compared to orders that increase (in absolute magnitude) a short position. This shows that all orders that initiate or increase a stock position are informed.

Similarly, in Table 8 Panel B, we focus on all orders that close or decrease an existing stock position (long sells and short buys) and test whether returns continue to drift in the opposite direction of these trades; i.e., whether post-order returns for long positions are still larger than those for short positions. The results confirm our earlier finding obtained for closing orders only (Table 3). Indeed, the magnitude of the return difference between long sells and short buys—about 0.6-1.6% over the 60-day horizon and 1.2-2.8% over the 125-day horizon—is even larger compared to what we found there. In other words, not only do hedge funds close their positions too early, they tend to reduce these positions even earlier, implying that more profits are foregone.

This robustness check confirms that our results are not driven by the fact that we focus exclusively on initial or closing orders, but that the results are robust to including follow-up orders.

VI. Conclusion

The question to which extent rational arbitrageurs remove mispricings and thus promote efficient prices is of crucial importance in asset pricing. Many arbitrageurs (and academics) focus on relative-value arbitrage; i.e., the process of exploiting relative price differences between two assets with comparable payoffs. However, a group of assets can be priced correctly relative to each other, while still being jointly mispriced when all assets are far away from their fundamental value. Thus, relative-value arbitrage in and of itself is not sufficient to ensure that prices reflect fundamental values. Similarly, while relative-value arbitrageurs arguably make prices more efficient (in the sense that they faster incorporate public information), they don't necessarily

make prices more informative (in the sense that prices summarize a higher absolute level of information).¹¹ In financial markets, this important void is filled by what we call *fundamental arbitrageurs*—traders that produce new and synthesize existing information in order to trade on deviations of asset prices from their fundamental values. Long-short equity hedge funds are presumably the most important group of investors belonging to this category. They spend substantial resources to gain an informational advantage and then take directional (long or short) bets in a relatively small number of stocks about which they have strong convictions.

In this paper, we exploit proprietary trading data for a sample of such long-short equity hedge funds to offer a microscopic analysis of their arbitrage activity. We first establish that positions initiated by these funds predict risk-adjusted returns over a horizon of up to one year, suggesting that their trades are informed. We then make the surprising observation that their closing trades are not only uninformed, but rather predict returns in the opposite direction of the closing trade. This implies that our sample hedge funds close positions that would have otherwise earned risk-adjusted returns going forward. In other words, they leave money on the table.

We then show that this behavior can be rationalized with the help of a simple trading model in which hedge funds are subject to a risk constraint and position monitoring costs and in which trading opportunities exhibit alpha decay. Under these assumptions, funds may rationally decide to close positions that are still expected to generate profits (1) in order to invest their limited capital in even more profitable trading opportunities or (2) in response to tightened capital or volatility constraints. We document supporting evidence for these predictions in the data.

Our findings have profound implications for our understanding of the limits of (fundamental) arbitrage. Indeed, we believe that we are the first to provide micro-level evidence on how rational arbitrageurs decide to abandon a profitable trading opportunity due to their risk capital and/or position monitoring constraints. As the trading opportunity is not fully exploited, mispricing persists. Thus, despite the presence of informed and rational arbitrageurs, market prices remain inefficient.

¹¹ See Brunnermeier (2005) and Weller (2016) on this point.

References

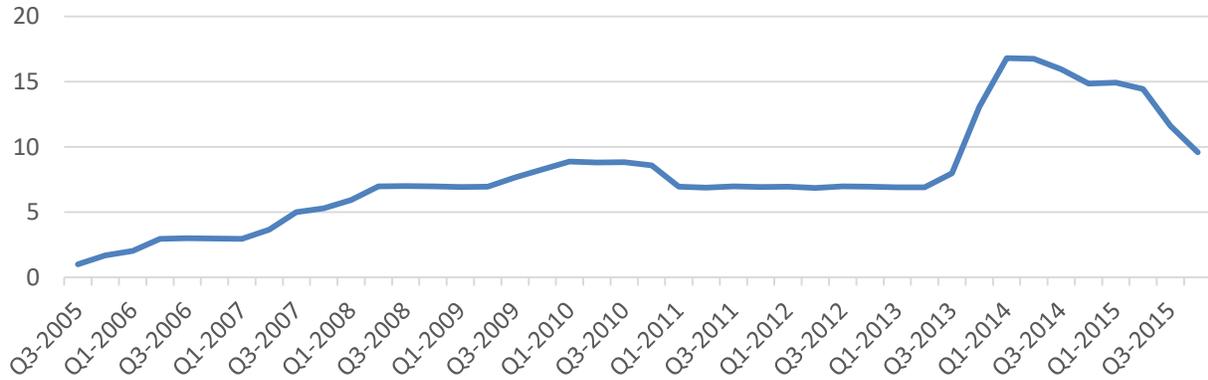
- Agarwal, Vikas, Kevin A. Mullally, and Narayan Y. Naik, 2015, Hedge Funds: A Survey of the Academic Literature, *Foundations and Trends in Finance*, forthcoming.
- Asquith, P., Pathak, P. A., and Ritter, J. R., 2005, Short interest, institutional ownership, and stock returns. *Journal of Financial Economics* 78, 24-276.
- Boehmer, Ekkehart, Truong X. Duong, and Zsuzsa R. Huszár, 2015, Short covering trades, *Working Paper*.
- Boehmer, Ekkehart, Charles M. Jones, and Xiaoyan Zhang, 2008, Which Shorts Are Informed?, *Journal of Finance* 63, 491-527.
- Brunnermeier, Markus K, 2005, Information Leakage and Market Efficiency, *Review of Financial Studies* 18, 417-457.
- Cao, Charles, Yong Chen, William N. Goetzmann, and Bing Liang, 2016, The role of hedge funds in the security price formation process, *Working paper*.
- Carhart, Mark M., 1997, On Persistence in Mutual Fund Performance, *The Journal of Finance* 52, 57-82.
- Choi, Jaewon, Neil D. Pearson, and Shastri Sandy, 2016, A First Glimpse into the Short Side of Hedge Funds, *Working Paper*.
- Daniel, Kent, Mark Grinblatt, Sheridan Titman and Russ Wermers, 1997, Measuring Mutual Fund Performance with Characteristic-Based Benchmarks, *Journal of Finance* 52, 1035-1058.
- Daniel, Kent, and Tobias J. Moskowitz, 2013, Momentum Crashes, *Working Paper*.
- Desai, H., Ramesh, K., Thiagarajan, S. R., Balachandran, B. V., 2002. Investigation of the information role of short interest in the NASDAQ Market. *Journal of Finance* 57, 2263-2287.
- Diether, Karl B., Kuan-Hui Lee, and Ingrid M. Werner, 2009, Short-Sale Strategies and Return Predictability, *Review of Financial Studies* 22, 575-607.
- Di Mascio, Rick, Anton Lines and Narayan Y. Naik, 2016, Alpha Decay and Strategic Trading, *Working Paper*.
- Engelberg, Joseph E., Adam V. Reed, and Matthew C. Ringgenberg, 2012, How are shorts informed?: Short sellers, news, and information processing, *Journal of Financial Economics* 105, 260-278.
- Foster, F. Douglas, and S. Viswanathan, 1996, Strategic Trading When Agents Forecast the Forecasts of Others, *The Journal of Finance* 51, 1437-1478.
- Franzoni, Francesco and Alberto Plazzi, 2015, What Constrains Liquidity Provision? Evidence From Hedge Fund Trades, *Working Paper*.
- Frazzini, Andrea and Lasse Heje Pedersen (2014). "Betting Against Beta", *Journal of Financial Economics* 111, 1-25.
- Griffin, John M. and Jun Xu, 2009, How smart are the smart guys? A unique view from hedge fund stock holdings, *Review of Financial Studies* 22, 2332-2370.
- Griffin, John M., 2002, Are the Fama and French Factors Global or Country Specific?, *Review of Financial Studies* 15, 783-803.
- Gromb, Denis, and Dimitri Vayanos, 2002, Equilibrium and welfare in markets with financially constrained arbitrageurs, *Journal of Financial Economics* 66, 361-407.
- Gromb, Denis, and Dimitri Vayanos, 2010, Limits of Arbitrage: The State of the Theory, *Annual Review of Financial Economics* 2, 251-275.
- Jame, Russell, 2016, Liquidity Provision and the Cross-Section of Hedge Fund Returns, *Working Paper*.

- Jank, Stephan, and Esad Smajlbegovic, 2015, Dissecting Short-Sale Performance: Evidence from Large Position Disclosures, *Working Paper*.
- Jylha, Petri, Kalle Rinne, and Matti Suominen, 2013, Do hedge funds supply or demand liquidity, *Review of Finance* 18, 1259-1298.
- Karolyi, Andrew and Ying Wu, 2014, Size, Value, and Momentum in International Stock Returns: A New Partial-Segmentation Approach, *Working Paper*.
- Kyle, Albert S., 1985, Continuous Auctions and Insider Trading, *Econometrica* 53, 1315-1336.
- Levi, Yaron and Ivo Welch, 2016, Assessing Cost-of-Capital Inputs, *Working Paper*.
- Pedersen, Lasse H., 2015, Efficiently Inefficient: How smart money invests & market prices are determined, *Princeton University Press*.
- Shleifer, Andrei, and Robert W. Vishny, 1997, The Limits of Arbitrage, *The Journal of Finance* 52, 35-55.
- Vasicek, Oldrich A. (1973). "A Note On Using Cross-Sectional Information In Bayesian Estimation Of Security Betas", *Journal of Finance* 28, 1233–1239.
- Weller, Brian M., 2016, Efficient Prices at Any Cost: Does Algorithmic Trading Deter Information Acquisition?, *Working Paper*.

Figure 1: Coverage over sample period

This figure shows the coverage over our sample period. Panel A shows the average number of funds in the sample for each quarter. Panel B shows the number of orders per day and of open positions per day averaged over the quarter.

Panel A: Number of funds in the sample



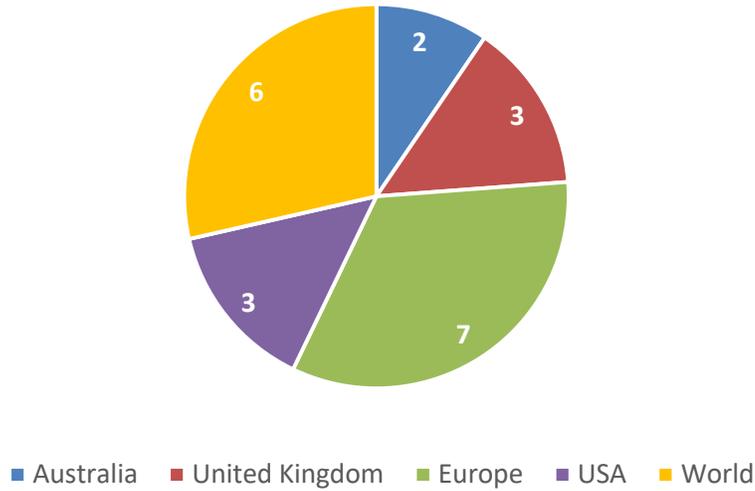
Panel B: Number of orders and positions per day



Figure 2: Investment areas of funds

Panel A shows the investment areas of our sample of funds. We base these areas on their chosen benchmark, but verify that the funds indeed invest predominantly in these regions. Panel B depicts the regions of the stocks held by the funds. We compute this average over the number of positions over the entire sample period. The definition of the regions are displayed in Appendix 2.

Panel A: Investment area of fund as specified by their benchmark



Panel B: Region of stocks held by funds (%)

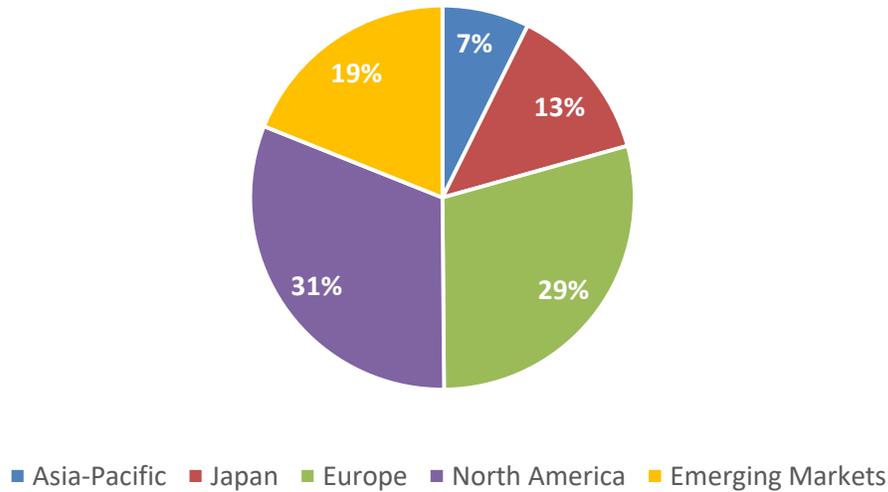
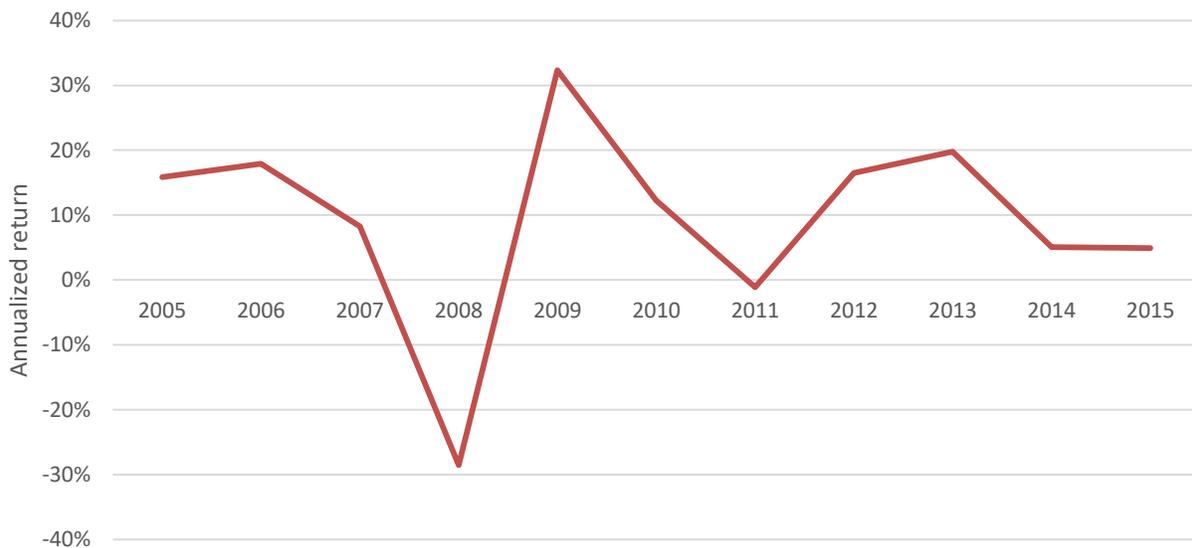


Figure 3: Fund profitability

In this figure, we display fund profitability by year. Fund profitability is computed from holdings data at the daily level and then annualized. In Panel A, we display profitability based on all long and short positions. In Panel B, we first compute profitability on the long and short side separately and then weight them equally.

Panel A: Fund profitability



Panel B: Fund profitability (long and short portfolio equally-weighted)

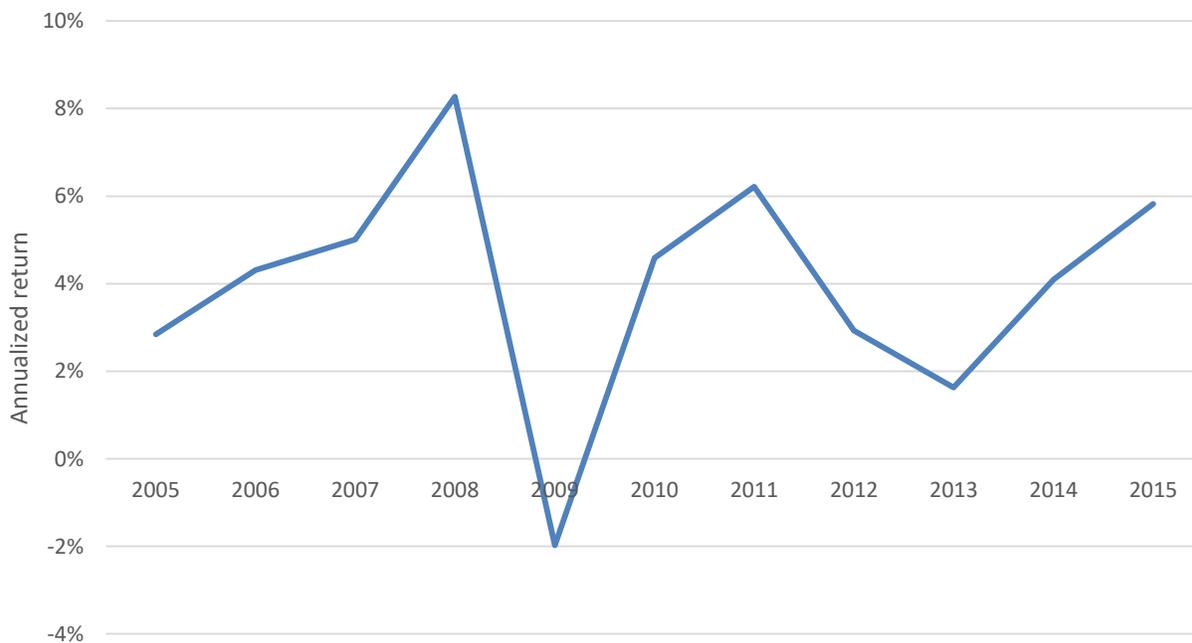


Figure 4: Excess returns following orders

This figure displays cumulative excess return indices for 250 trading days following orders that initiate or close a position. *Initiate Long Position* is the buy order establishing a long position. *Initiate Short Position* is the sell order establishing a short position. *Close Short Position* is the buy order closing a short position. *Close Long Position* is the sell order closing a long position. *Excess return* is the return of the stock minus the return of the fund-specified benchmark. The return index is set to 100 at the last day of the order. To reduce the effect of outliers, we set daily returns above 10% to 10% and below -10% to -10%.

Excess returns around orders

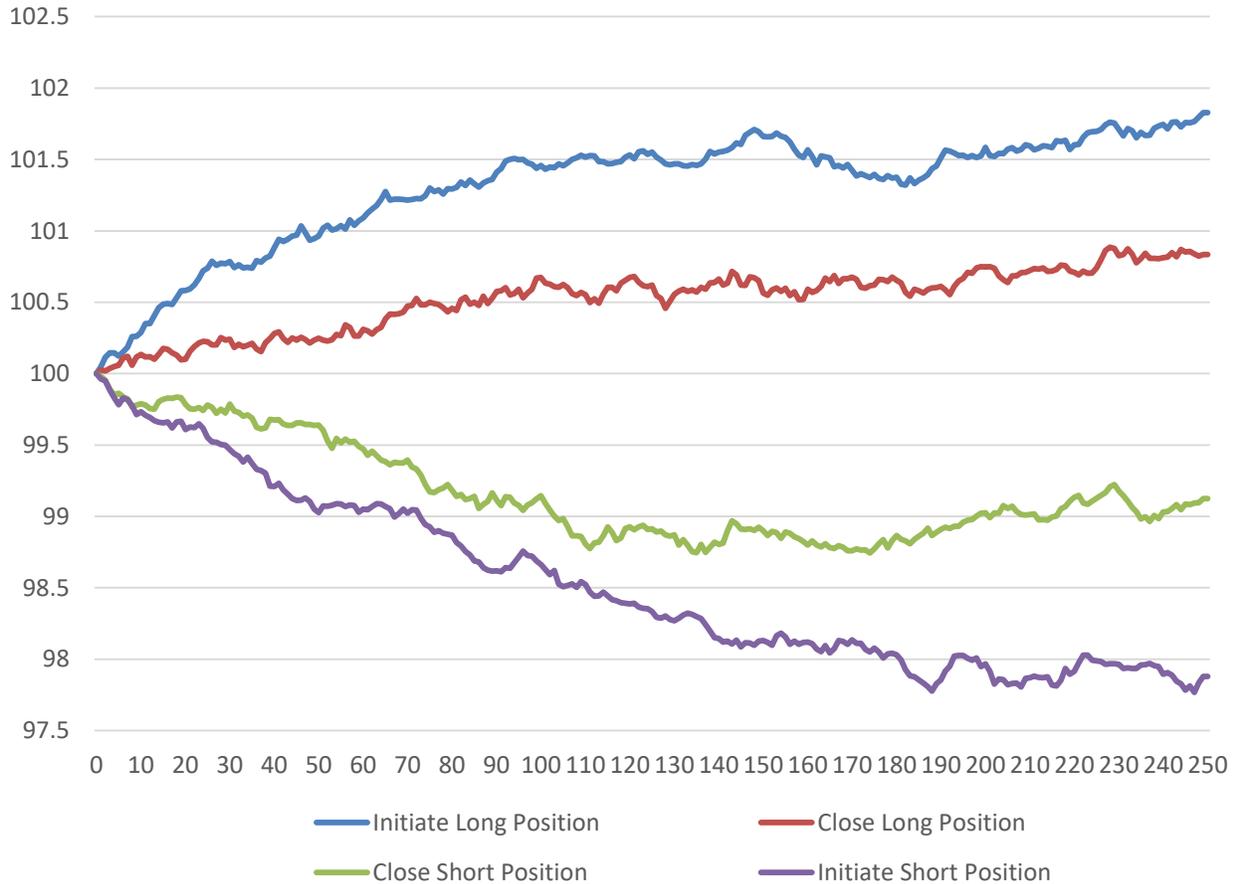


Table 1: Summary statistics

Panel A displays summary statistics by fund. *Number of Long (Short) Positions* is the average of long (short) positions held by the fund. *Short Fraction* is the average fraction that short positions constitute of the fund holdings (measured in USD). *Orders per Day* are the average number of orders initiated per day. *Trade Fraction* is the average of the funds trading volume divided by the value of its holdings. *Value of Fund* is the average value of all positions (long and short combined). *Positions per Stock* is the average number of times the fund establishes a position in the specific stock. Panel B displays summary statistics by position. A position lasts from its initiation (first buy for long positions or first sell for short positions) to its close (moment the holding of the stock is reduced to zero). *Number of Direction Changes* is the number of times the orders move from buy to sell orders or from sell to buy orders while the position is open. *Open Start* is a dummy variable equal to one if the position is open already at the time the fund enters the database. *Open End* is a dummy variable equal to one if the position is still open when the fund leaves the database. Panel C-E display summary statistics by order. We split the orders by whether they initiate a position, close a position or simply change the size of a position. *Size as fraction of largest holding* is the size of the order in terms of number of stocks traded divided by the largest size the position ever has.

Panel A: Averages by fund

Variable	Mean	10 th Percentile	Median	90 th Percentile	Standard Deviation
Number of Long Positions	49.8	16.9	36.1	74.9	43.4
Number of Short Positions	23.9	10.8	18.6	46.3	14.2
Short Fraction (%)	30.2	15.8	26.4	48.7	19.2
Orders per Day	5.81	1.54	5.60	10.5	3.58
Trade Fraction (%)	5.37	0.82	2.75	14.0	5.36
Value of Fund (billion USD)	2.05	0.12	0.35	6.41	3.63
Positions per Stock	1.96	1.37	1.90	2.70	0.61
Observations	21				

Panel B: Statistics by position

Variable	Mean	10 th Percentile	Median	90 th Percentile	Standard Deviation
Length (trading days)	104.4	4	35	275	188.9
Number of Orders	5.92	2	3	12	8.89
Number of Direction Changes	2.50	1	1	5	5.06
Open Start	0.069	0	0	0	0.25
Open End	0.11	0	0	1	0.32
Observations	16241				

Panel C: Statistics by order – initial orders

Variable	Mean	10 th Percentile	Median	90 th Percentile	Standard Deviation
Number of Trades	1.61	1	1	3	1.54
USD volume (million USD)	11.3	0.27	3.64	22.5	40.8
Size as fraction of largest holding (%)	77.3	24.8	100	100	31.0
Observations	15114				

Panel D: Statistics by order – follow-up orders

Variable	Mean	10 th Percentile	Median	90 th Percentile	Standard Deviation
Number of Trades	1.49	1	1	3	1.30
USD volume (million USD)	7.66	0.085	1.70	17.1	31.5
Size as fraction of largest holding (%)	15.8	0.92	8.65	42.6	18.3
Observations	66700				

Panel E: Statistics by order – closing orders

Variable	Mean	10 th Percentile	Median	90 th Percentile	Standard Deviation
Number of Trades	1.64	1	1	3	1.94
USD volume (million USD)	10.9	0.23	3.30	22.2	34.4
Size as fraction of largest holding (%)	78.2	25.7	100	100	31.0
Observations	14330				

Table 2: Returns following the initiation of positions

This table examines returns following the initiation of positions. We include only orders that initiate a position. We regress average returns following the order on a dummy variable whether the order is related to a short position (and is thus a short sell). The dependent variable is the cumulative return expressed in percent for 60 and 125 trading days following the last day of the order. *Excess return* is the return of the stock minus the return of the fund-specified benchmark. *DGTW return* is the return of the stock minus the average return of a portfolio sorted by region, size, book-to-market and momentum. *Four-factor Alpha* is the alpha according to the Carhart (1997) model estimated at the regional level. Details on variable constructions can be found in Appendix 1. Month fixed effects are based on the month of the last day of the order. All standard errors are two-way clustered by stock and last date of order. We report t-statistics below the coefficients in parenthesis. ***, **, * indicate significance at the 1%, 5% and 10% level.

Dependent Variable:	Excess Return t+1, t+60	Excess Return t+1, t+125	DGTW Return t+1, t+60	DGTW Return t+1, t+125	4-Factor Alpha t+1, t+60	4-Factor Alpha t+1, t+125
	(1)	(2)	(3)	(4)	(5)	(6)
D(Short)	-1.86*** (-5.95)	-2.60*** (-5.47)	-1.59*** (-5.62)	-1.95*** (-4.68)	-1.66*** (-5.61)	-2.04*** (-4.67)
Observations	13050	12409	11700	11214	12316	11898
Adjusted R ²	0.06	0.09	0.04	0.05	0.02	0.02
Fund Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Table 3: Returns following the closure of positions

This table examines returns following the closure of positions. We include only orders that close a position. We regress average returns following the order on a dummy variable whether the order is related to a short position (and is thus a short sell). The dependent variable is the cumulative return expressed in percent for 60 and 125 trading days following the last day of the order. *Excess return* is the return of the stock minus the return of the fund-specified benchmark. *DGTW return* is the return of the stock minus the average return of a portfolio sorted by region, size, book-to-market and momentum. *Four-factor Alpha* is the alpha according to the Carhart (1997) model estimated at the regional level. Details on variable constructions can be found in Appendix 1. Month fixed effects are based on the month of the last day of the order. All standard errors are two-way clustered by stock and last date of order. We report t-statistics below the coefficients in parenthesis. ***, **, * indicate significance at the 1%, 5% and 10% level.

Dependent Variable:	Excess Return t+1, t+60	Excess Return t+1, t+125	DGTW Return t+1, t+60	DGTW Return t+1, t+125	4-Factor Alpha t+1, t+60	4-Factor Alpha t+1, t+125
	(1)	(2)	(3)	(4)	(5)	(6)
D(Short)	-0.73** (-2.29)	-1.50*** (-3.05)	-0.30 (-0.99)	-1.03** (-2.35)	-0.37 (-1.16)	-0.93** (-2.01)
Observations	11952	11320	11027	10469	11598	11114
Adjusted R ²	0.08	0.11	0.04	0.06	0.02	0.03
Fund Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Table 4: Comparing returns following initiation and closing orders

This table compares returns following the initiation and closing of positions. We regress returns following the order on a dummy variable equal to one if the trade is an initiation order. We only include initiation and closing orders. In Panel A, we only include orders related to long positions. In Panel B, we only include orders related to short positions. In Panel C, we include orders related to both long and short positions. In Panel C, the dependent variables are “profits” based on different return measures. Which are returns that are multiplied by negative one for short positions. The dependent variable is the cumulative return expressed in percent for 60 and 125 trading days following the last day of the order. Details on variable constructions can be found in Appendix 1. Month fixed effects are based on the month of the last day of the order. All standard errors are two-way clustered by stock and last date of order. We report t-statistics below the coefficients in parenthesis. ***, **, * indicate significance at the 1%, 5% and 10% level.

Panel A: Long Positions

Dependent Variable:	Excess Return t+1, t+60	Excess Return t+1, t+125	DGTW Return t+1, t+60	DGTW Return t+1, t+125	4-Factor Alpha t+1, t+60	4-Factor Alpha t+1, t+125
	(1)	(2)	(3)	(4)	(5)	(6)
D(Initiation)	0.76*** (3.21)	0.96*** (2.82)	0.86*** (4.14)	0.67** (2.44)	0.64*** (2.92)	0.79*** (2.73)
Observations	13419	12779	11980	11457	12743	12299
Adjusted R ²	0.06	0.09	0.04	0.05	0.02	0.03
Fund Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Short Positions

Dependent Variable:	Excess Return t+1, t+60	Excess Return t+1, t+125	DGTW Return t+1, t+60	DGTW Return t+1, t+125	4-Factor Alpha t+1, t+60	4-Factor Alpha t+1, t+125
	(1)	(2)	(3)	(4)	(5)	(6)
D(Initiation)	-0.34 (-1.32)	-0.31 (-0.91)	-0.53** (-2.30)	-0.37 (-1.23)	-0.64** (-2.37)	-0.45 (-1.31)
Observations	11583	10950	10747	10226	11171	10713
Adjusted R ²	0.08	0.11	0.04	0.05	0.03	0.02
Fund Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Panel C: Long and Short Positions

Dependent Variable:	Signed Excess Return t+1, t+60	Signed Excess Return t+1, t+125	Signed DGTW Return t+1, t+60	Signed DGTW Return t+1, t+125	Signed 4-Factor Alpha t+1, t+60	Signed 4-Factor Alpha t+1, t+125
	(1)	(2)	(3)	(4)	(5)	(6)
D(Initiation)	0.60*** (3.15)	0.73*** (2.62)	0.73*** (4.56)	0.56** (2.47)	0.67*** (3.87)	0.65*** (2.82)
Observations	25002	23729	22727	21683	23914	23012
Adjusted R ²	0.00	0.00	0.00	0.00	0.00	0.00
Fund Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Table 5: Returns following the closure of positions – Split by Position Changes

This table examines whether returns following the closure of positions dependent on changes in the number of positions of the fund. We run the same regression as in Table 3 but split the sample by whether the fund increased or decreased the number of positions before the first day of the order. In Panel A, we split by change in number of positions in the 5 days prior to the order. In Panel B, we split by change in number of positions in the 10 days prior to the order. We include only orders that close a position. We regress average returns following the order on a dummy variable whether the order is related to a short position (and is thus a short sell). The dependent variable is the cumulative return expressed in percent for 60 and 125 trading days following the last day of the order. Details on variable constructions can be found in Appendix 1. Month fixed effects are based on the month of the last day of the order. All standard errors are two-way clustered by stock and last date of order. We report t-statistics below the coefficients in parenthesis. ***, **, * indicate significance at the 1%, 5% and 10% level.

Panel A: Split by change in number of positions relative to 5 trading days before

Dependent Variable:	Excess Return t+1, t+125		DGTW Return t+1, t+125		4-Factor Alpha t+1, t+125	
Sample	More Positions	Less Positions	More Positions	Less Positions	More Positions	Less Positions
	(1)	(2)	(3)	(4)	(5)	(6)
D(Short)	-2.51*** (-3.35)	-0.62 (-1.01)	-1.61** (-2.54)	-0.54 (-0.91)	-2.02*** (-2.95)	-0.09 (-0.14)
Observations	5080	6227	4716	5743	4985	6118
Adjusted R ²	0.11	0.12	0.06	0.06	0.03	0.03
Fund Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Split by change in number of positions relative to 10 trading days before

Dependent Variable:	Excess Return t+1, t+125		DGTW Return t+1, t+125		4-Factor Alpha t+1, t+125	
Sample	More Positions	Less Positions	More Positions	Less Positions	More Positions	Less Positions
	(1)	(2)	(3)	(4)	(5)	(6)
D(Short)	-2.00*** (-2.80)	-1.31** (-2.03)	-1.36** (-2.34)	-0.97 (-1.60)	-1.75*** (-2.75)	-0.42 (-0.66)
Observations	5478	5829	5098	5361	5373	5730
Adjusted R ²	0.11	0.12	0.07	0.04	0.04	0.02
Fund Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Table 6: Returns following the closure of positions – Split by Fund Profits

This table examines whether returns following the closure of positions dependent on the profitability of the fund before the order. We run the same regression as in Table 3 but split the sample by whether the fund had positive or negative returns before the first day of the order. In Panel A, we split by fund return in the 5 days prior to the order. In Panel B, we split by fund return in the 10 days prior to an order. We include only orders that close a position. We regress average returns following the order on a dummy variable whether the order is related to a short position (and is thus a short sell). The dependent variable is the cumulative return expressed in percent for 60 and 125 trading days following the last day of the order. Details on variable constructions can be found in Appendix 1. Month fixed effects are based on the month of the last day of the order. All standard errors are two-way clustered by stock and last date of order. We report t-statistics below the coefficients in parenthesis. ***, **, * indicate significance at the 1%, 5% and 10% level.

Panel A: Split by fund return over prior 5 trading days

Dependent Variable:	Excess Return t+1, t+125		DGTW Return t+1, t+125		4-Factor Alpha t+1, t+125	
Sample	Positive Fund Return	Negative Fund Return	Positive Fund Return	Negative Fund Return	Positive Fund Return	Negative Fund Return
	(1)	(2)	(3)	(4)	(5)	(6)
D(Short)	-0.43 (-0.75)	-2.80*** (-3.67)	-0.74 (-1.34)	-1.57** (-2.35)	-0.37 (-0.65)	-1.68** (-2.35)
Observations	6350	4970	5859	4610	6222	4892
Adjusted R ²	0.10	0.13	0.05	0.07	0.03	0.04
Fund Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Split by fund return over prior 10 trading days

Dependent Variable:	Excess Return t+1, t+125		DGTW Return t+1, t+125		4-Factor Alpha t+1, t+125	
Sample	Positive Fund Return	Negative Fund Return	Positive Fund Return	Negative Fund Return	Positive Fund Return	Negative Fund Return
	(1)	(2)	(3)	(4)	(5)	(6)
D(Short)	-0.56 (-0.88)	-2.74*** (-3.69)	-0.43 (-0.79)	-1.69*** (-2.63)	-0.27 (-0.44)	-1.91*** (-2.75)
Observations	6562	4758	6065	4404	6442	4672
Adjusted R ²	0.10	0.13	0.06	0.07	0.03	0.03
Fund Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Table 7: Returns following the closure of positions – Split by Change in Fund Volatility

This table examines whether returns following the closure of positions dependent on changes in fund volatility. We run the same regression as in Table 3 but split the sample by whether the fund experienced an increase or a decrease in its volatility. Our measure of fund volatility is the standard deviation of daily fund returns over a month. We compare this volatility in the month of the order relative to the volatility in the prior month. We include only orders that close a position. We regress average returns following the order on a dummy variable whether the order is related to a short position (and is thus a short sell). The dependent variable is the cumulative return expressed in percent for 60 and 125 trading days following the last day of the order. Details on variable constructions can be found in Appendix 1. Month fixed effects are based on the month of the last day of the order. All standard errors are two-way clustered by stock and last date of order. We report t-statistics below the coefficients in parenthesis. ***, **, * indicate significance at the 1%, 5% and 10% level.

Dependent Variable:	Excess Return t+1, t+125		DGTW Return t+1, t+125		4-Factor Alpha t+1, t+125	
Sample	Higher Volatility	Lower Volatility	Higher Volatility	Lower Volatility	Higher Volatility	Lower Volatility
	(1)	(2)	(3)	(4)	(5)	(6)
D(Short)	-2.04*** (-3.17)	-0.90 (-1.22)	-1.65*** (-2.70)	-0.37 (-0.60)	-0.90 (-1.41)	-0.89 (-1.35)
Observations	5814	5282	5360	4896	5722	5173
Adjusted R ²	0.10	0.14	0.07	0.07	0.04	0.04
Fund Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Table 8: Robustness checks including all orders

This table shows robustness checks for tables 2 and 3. In Panel A, we show the same analysis as in Table 2 but we include all orders that increase a position instead of all initiating orders, i.e. we include all long buys and short sells. In Panel B, we show the same analysis as in Table 3 but we include all orders that decrease a position instead of all closing orders. Details on variable constructions can be found in Appendix 1. Month fixed effects are based on the month of the last day of the order. All standard errors are two-way clustered by stock and last date of order. We report t-statistics below the coefficients in parenthesis. ***, **, * indicate significance at the 1%, 5% and 10% level.

Panel A: Robustness to Table 2 – Orders increasing a position

Dependent Variable:	Excess Return t+1, t+60	Excess Return t+1, t+125	DGTW Return t+1, t+60	DGTW Return t+1, t+125	4-Factor Alpha t+1, t+60	4-Factor Alpha t+1, t+125
	(1)	(2)	(3)	(4)	(5)	(6)
D(Short)	-2.04*** (-5.08)	-3.08*** (-5.41)	-1.30*** (-4.95)	-2.08*** (-4.82)	-1.16*** (-3.98)	-1.58*** (-3.19)
Observations	43444	41178	38935	37248	41592	39938
Adjusted R ²	0.02	0.04	0.03	0.04	0.02	0.02
Fund and Month F.E.	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Robustness to Table 3 – Orders decreasing a position

Dependent Variable:	Excess Return t+1, t+60	Excess Return t+1, t+125	DGTW Return t+1, t+60	DGTW Return t+1, t+125	4-Factor Alpha t+1, t+60	4-Factor Alpha t+1, t+125
	(1)	(2)	(3)	(4)	(5)	(6)
D(Short)	-1.59*** (-4.41)	-2.79*** (-4.88)	-0.90*** (-3.30)	-1.57*** (-3.45)	-0.57* (-1.90)	-1.24** (-2.37)
Observations	40688	38205	36583	34578	39306	37247
Adjusted R ²	0.02	0.05	0.04	0.06	0.02	0.03
Fund and Month F.E.	Yes	Yes	Yes	Yes	Yes	Yes

Appendix A.1: Variable definitions

This table displays the variable definitions for all variables used in the regressions.

Variable Name	Definition
Excess Return	$Stock\ Return - Benchmark\ Return$
Stock Return	Return in USD from Datastream.
Benchmark Return	USD return of the benchmark specified by the fund. The benchmark is specific for the fund, but is the same for both long and short positions of the fund. Data is provided by Analytics.
DGTW Return	$Stock\ Return - Return\ of\ portfolio\ of\ similar\ stocks$ Similar stocks are stocks in the same quintile of market capitalization, book-to-market ratio and 12 month momentum within the same region. For more details see Appendix 3.
4-Factor Alpha	$r_{c,t} - \beta_c * (r_{m,t} - r_{f,t}) - \beta_{HML} * HML_t - \beta_{SMB} * SMB_t - \beta_{MOM} * MOM_t$ For more details see Appendix 3.
Signed Excess Return	Excess Return for long positions. $Excess\ Return * -1$ for short positions.
Signed DGTW Return	DGTW Return for long positions. $DGTW\ Return * -1$ for short positions.
Signed 4-Factor Alpha	4-Factor Alpha for long positions. $4\ Factor\ Alpha * -1$ for short positions.
D(Short)	Dummy variable equal to one if the order related to a short position and zero if it is related to a long position.
D(Buy)	Dummy variable equal to one if the order is a buy order and zero if the order is a sell order.
D(Buy) - Holding	Dummy variable equal to one if the day contained a buy trade and zero if there was no trade or a sell trade. Notice that this variable is constructed based on trades rather than orders, so the variable is equal to 1 on the day of the second trade of a buy order.
D(Sell) - Holding	Dummy variable equal to one if the day contained a sell trade and zero if there was no trade or a buy trade. Notice that this variable is constructed based on trades rather than orders, so the variable is equal to 1 on the day of the second trade of a sell order.
Direction	Variable equal to one if the day contained a buy trade, minus one if it contained a sell trade and zero if there was no trade. Notice that this variable is constructed based on trades rather than orders, so the variable is equal to 1 on the day of the second trade of a buy order.
Position Tenure	$\log(Number\ of\ days\ the\ position\ is\ open)$
Fund profitability	Value-weighted average return of all positions of the fund, where the weight is the dollar value of the position. The returns of short positions are multiplied by minus one.
Fund volatility	Monthly standard deviation of daily fund returns. Volatility is set to missing when we have fewer than 15 non-missing daily return observations for a given month.

Appendix A.2: Regions

This table displays the regions to which the countries are assigned. The region assignments follows Karolyi and Wu (2014)

Panel A: Region assignments

Country Name	Region
Japan	Japan
Canada	North America
United States	North America
Australia	Asia-Pacific
New Zealand	Asia-Pacific
Singapore	Asia-Pacific
Hong Kong	Asia-Pacific
Austria	Europe
Belgium	Europe
Denmark	Europe
Finland	Europe
France	Europe
Germany	Europe
Greece	Europe
Ireland	Europe
Italy	Europe
Netherlands	Europe
Norway	Europe
Portugal	Europe
Spain	Europe
Sweden	Europe
Switzerland	Europe
United Kingdom	Europe
Argentina	EME
Brazil	EME
Chile	EME
China	EME
Colombia	EME
Czech Republic	EME
Hungary	EME
India	EME
Indonesia	EME
Israel	EME
Korea (South)	EME
Malaysia	EME
Mexico	EME
Pakistan	EME
Peru	EME
Philippines	EME
Poland	EME
Russian Federation	EME
South Africa	EME
Taiwan	EME
Thailand	EME
Turkey	EME
Venezuela	EME

Appendix A.3: Additional Information on Dataset Construction

1) Merging of datasets

We merge the trading and the holding dataset provided by Inalytics. We first merge based on ISIN. Trades that we cannot match by ISIN, we match by SEDOL and finally by CUSIP. Theoretically, whenever there is a change in the number of shares held in the holdings data, we would expect to see a corresponding trade in the trade data. In fact, there are some errors in the data and the trade and holding data do not match perfectly. According to Inalytics, the holding data are more accurate. Therefore, we rely on the holdings data, i.e. we assume there is a trade whenever there is a change of holding in the holdings data. The only exception is if a holding disappears from the data and then reappears shortly afterwards. In these cases we fill in the missing dates in between with the old quantity. Reappearing shortly afterwards means within 22 trading days (one month); or within 70 trading days (one quarter) if the position reappears with the exact same number of stocks. In total this fill-in procedure affects only about 3400 stock-days in the holding data (compared to about 1.7 million stock-days in the full sample).

In total, we have about 150,000 trades according to the holdings data. For about 90% of these trades, we have a corresponding trade in the trading data. However, for only about 83% of these trades does the number of stocks traded according to the trading data match the change in the number of stocks held in the holding data. Following Inalytics' advice, we assume the holdings data to be correct, but still use the trading price from the trading data even if the number of stocks traded does not match. In addition, we set unrealistic trade prices to missing. We determine if a trade price is realistic by comparing it to the closing prices around the trade. If the trade price is more than 5% outside of the bound between these two prices, we set it to missing. This correction affects about 2000 trades. In this paper, we only use trade prices to compute summary statistics.

2) Stock Universe

To compute DGTW returns and regional factors, we need a universe of stocks. We construct this stock universe by matching Worldscope and Datastream data. We only keep stocks that are covered in both databases. We only keep one stock per company (we identify companies using Worldscope Permanent Identifier). We also only keep stocks from the countries listed in Appendix 2 (we take the country from Worldscope). We require stock to have a positive book value, information on market capitalization in Worldscope and a stock price of at least USD 0.20.

3) Stock Returns and Balance Sheet Data

We download daily returns for stocks in our stock universe and in the Inalytics data from Datastream using ISINs (and then using SEDOLs if we do not find a match using ISINs). We download returns in local currency and convert them to USD using the exchange rates on Datastream. Using local currency returns minimizes

the errors due to rounding for stocks with low stock prices. When computing returns for the DGTW portfolios and Carhart (1997) factor, we remove returns in the top and bottom 0.25% by region following the instructions on the website of Kenneth French. We compare the regional factors that we compute to those provided by Kenneth French on his website and they are very similar. Because we need factors for the emerging market region and French does not provide them, we have to compute factors ourselves at least for this region. To keep everything comparable across regions, we use our self-computed factors for other regions as well.

We take market capitalization in USD directly from Worldscope (code 07210) and compute book-to-market directly from Worldscope as the inverse of the price-to-book ratio (code 09304). We use annual Worldscope data.

For stocks in the Analytics data that are not covered in Datastream, we receive stock return information from Analytics. Because we don't have balance sheet information for these stocks, we cannot compute DGTW returns. But we can compute excess returns and alphas. Of about 1.7 million stock-days in which a position is open, we have 1.43 million (84%) observations with return data on Datastream. By filling in the Analytics return data we can increase this coverage to 1.66 million stock-days with returns (98%).

4) Four-Factor alphas

To compute alphas, we construct daily four-factors of the Carhart (1997) model for each of our 5 regions (see Appendix 2) following the instructions on the website of Kenneth French¹². We use the U.S. one month T-bill rate as the risk free rate and all returns are in U.S. dollars. We compute market returns as value-weight average returns for our stock universe in each region (the stock universe construction is describe above). To construct the SMB and HML factors, we sort stocks in a region into two market cap and three book-to-market equity (B/M) groups at the end of each June. Big stocks are those in the top 90% of market cap for the region, and small stocks are those in the bottom 10%. The B/M breakpoints for a region are the 30th and 70th percentiles of B/M for the big stocks of the region. For the 6 portfolios thus formed, we compute value weighted returns for each day and then compute the factors as:

$$SMB = \frac{1}{3} * (Small Value + Small Neutral + Small Growth) - \frac{1}{3} * (Big Value + Big Neutral + Big Growth)$$

$$HML = \frac{1}{2} * (Small Value + Big Value) - \frac{1}{2} * (Small Growth + Big Growth)$$

The 2x3 sorts on size and lagged momentum to construct MOM are formed monthly. For portfolios formed at the end of month t-1, the lagged momentum return is a stock's cumulative return for month t-12 to month

¹² http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. I quote directly from these instructions in the description below. At the time of the construction of these factors, daily factors were not yet available on Ken French's website.

t-2. The momentum breakpoints for a region are the 30th and 70th percentiles of the lagged momentum returns of the big stocks of the region. For the 6 portfolios thus formed, I compute value weighted returns for each day and then compute the momentum factor as:

$$MOM = \frac{1}{2} * (Small\ High + Big\ High) - \frac{1}{2} * (Small\ Low + Big\ Low)$$

For each stock and each month, we then compute the beta with respect to those factors from a daily regression over the past year. We remove returns from the regression that are in the top and bottom 0.25% by region. Furthermore, we only keep betas that are based on at least 50 days of data.

$$r_{c,t} = \alpha + \beta_m * (r_{m,t} - r_{f,t}) + \beta_{HML} * HML + \beta_{SMB} * SMB + \beta_{MOM} * MOM$$

Where $r_{c,t}$ is the daily company return, $r_{m,t}$ is the daily market return and $r_{f,t}$ is the daily risk free rate.

Following Frazzini and Pedersen (2014), I shrink the betas to their cross sectional mean by computing:

$$\beta_{c,t}^{shrunk} = 0.7 * \beta_{c,t} + 0.3 * \beta_{region,t}$$

Then we use these shrunk betas to compute daily alphas as follows:

$$Four\ factor\ alpha_{c,t} = r_{c,t} - \beta_c * (r_{m,t} - r_{f,t}) - \beta_{HML} * HML_t - \beta_{SMB} * SMB_t - \beta_{MOM} * MOM_t$$

We compute alphas for all stocks in our region even if these stocks do not match the criteria to be included in the stock universe.

5) DGTW returns

To compute DGTW returns, we split the stocks in our universe (the stock universe construction is describe above) into 625 portfolios. First, we split the universe into the 5 geographic region (see Appendix 2). Second, each year, within each region we sort stocks into 5 portfolios by market capitalization. Third, within each of these 25 size-region portfolios we sort stocks by book-to-market. Fourth, within each of these 125 regio-size-book to market portfolios, we sort stocks into 5 portfolios by returns over months t-12 to t-2. While splits for market cap and market-to-book happen once a year, splits by past return are executed every month.

Then we compute the average return within each of the 625 portfolio as the average USD return in the portfolio weighted by market cap. Finally, we compute DGTW returns as stock return minus the return of the respective portfolio. We only compute a DGTW return if the stock was part of the portfolio.

Appendix B: A simple Hedge Fund Trading Model

1) General Setup

Time is discrete. In a given period t , a hedge fund has W_t units of capital and faces different investment opportunities: it can invest in a riskless asset with a net return of $r_f = 0$ or into N risky stocks that are all assumed to be uncorrelated with each other (and over time). Any stock i is either fairly priced or mispriced. If it is fairly priced, then its return from one period to the next is given by $r_{it} = \varepsilon_{it}$ where ε_{it} is a zero-mean noise term with variance σ_t^2 (constant for all i).¹³ If stock i is mispriced, then its return is given by $r_{it} = \Delta_{it} + \varepsilon_{it}$ where $\Delta_{it} > 0$ ($\Delta_{it} < 0$) captures the underpricing (overpricing). To capture the empirical reality that such trading opportunities disappear over time, we assume that the mispricing Δ_{it} decays over subsequent periods. That is, if the mispricing of stock i occurs in period t and lasts for τ periods, then we have $|\Delta_{it}| > |\Delta_{it+1}| > \dots > |\Delta_{it+\tau}| = 0$. The hedge fund is assumed to know which stocks are mispriced and by how much.

The stock return component ε_{it} represents fluctuations in the stock's fair value driven by public news. There are two ways to think about the mispricing in this setup: First, it could be that, occasionally, a stock price movement ε_{it} occurs that is not justified by fundamentals. After such an occurrence, the hedge fund learns about this mispricing and expects it to revert over time.¹⁴ Second, it could be that the hedge fund obtained private information that some future dividend is going to be higher/lower than expected, and the hedge fund expects this information to leak to the market over time.¹⁵

We assume that the hedge fund maximizes expected returns (i) after accounting for a *position-monitoring* cost and (ii) subject to not exceeding a volatility limit $\bar{\sigma}$. Let $\mathbf{w}_t = (w_{1t} \ \dots \ w_{Nt})'$ denote the vector of portfolio weights of the N risky stocks, w_{ft} be the portfolio weight of the riskfree asset, and $E(\mathbf{r}_t) = (E(r_{1t}) \ \dots \ E(r_{Nt}))'$ be the vector of expected stock returns (where $E(r_{it}) = \Delta_{it}$ if stock i is mispriced and zero otherwise). Furthermore, let $\mathbb{1}_{w_{it} \neq 0}$ be a dummy variable that takes the value one if the portfolio weight of stock i is strictly positive and zero otherwise. Let $\mathbb{1}_{\mathbf{w}_t}$ be the N -dimensional vector of these dummy variables.

Formally, the hedge fund's objective in period t is given by:

$$\max_{\mathbf{w}_t, w_{ft}, N_{Pt}} W_t [1 + \mathbf{w}_t' E(\mathbf{r}_{t+1})] - c N_{Pt} \quad \text{subject to}$$

$$\mathbf{w}_t' \mathbf{1}_N + w_{ft} = 1 \quad (1)$$

$$\mathbf{w}_t' \mathbf{w}_t \sigma_t^2 \leq \bar{\sigma}^2 \quad (2)$$

$$\mathbb{1}_{\mathbf{w}_t}' \mathbf{1}_N = N_{Pt} \quad (3)$$

¹³ Because all stocks are uncorrelated, there is no systematic risk in the economy and hence the risk premium is zero. Thus, the riskfree asset can be viewed as an investment in the perfectly-diversified market portfolio. Note that the assumption of a zero risk premium is without loss of generality: we obtain the same predictions if we assume that the hedge fund chases mispricings as a source of α ; i.e., any return in excess of the market risk premium.

¹⁴ Thus, in the example where the hedge fund becomes aware of the mispricing at the beginning of period t and expects it to last for τ periods, the return of stock i in period $t-l$ was given by $\varepsilon_{it-l} = \prod_{\pi=0}^{l-1} (1 + \Delta_{t+\pi})^{-1} - 1$ and was not justified by fundamental news.

¹⁵ For instance, the hedge fund could have learnt that the dividend in period $t+\tau-l$ is higher than expected by an amount which would justify a compounded return of $\prod_{\pi=0}^{l-1} (1 + \Delta_{t+\pi}) - 1$.

where $\mathbf{1}_N$ denotes the N -dimensional unit vector, N_{p_t} denotes the total number of open stock positions (i.e., positions with $w_{it} \neq 0$), and $c > 0$ is the monitoring cost *per* open position. Constraint (1) describes the standard *portfolio additivity* condition that portfolio weights have to sum to one. Constraint (2) ensures that the volatility of the entire portfolio is less or equal to $\bar{\sigma}$. This *risk limit* constraint is supposed to reflect risk management practices common among hedge funds (such as risk-parity investment). Constraint (3) says that the number of stock positions with non-zero portfolio weights equals N_{p_t} , which then causes total position monitoring costs of cN_{p_t} . This *position monitoring cost* is supposed to reflect the fact that the monitoring and management of directional equity bets requires a significant amount of limited attention (which can be relaxed by acquiring additional attention capacity at the cost c). As a result of this assumption, the hedge fund may choose to invest in fewer than the total number of mispriced stocks. Empirically, it is well-known that discretionary long-short hedge funds—in contrast to most institutional investors that seek well-diversified portfolios—choose to hold a fairly limited number of open positions.

It is important to understand that without the risk limit constraint, the hedge fund’s trading strategy and hence its profits would become unbounded. Indeed, by short-selling any overpriced stock (or the riskfree asset) and levering-up its positions in underpriced stocks, the hedge fund could increase profits without violating the portfolio additivity constraint. However, since such a strategy also increases the portfolio’s risk, the risk limit constraint prevents this case from occurring. This argument makes clear that the optimal trading rule must be such that the risk limit is exactly binding (as long as there is at least one mispriced stock).

In our simple setup, the hedge fund is myopic in that it maximizes its expected wealth one period ahead: $E(W_{t+1}) = W_t(1 + E(r_{p_{t+1}})) - cN_{p_t}$, where $r_{p_{t+1}} \equiv \mathbf{w}'_t \mathbf{r}_{t+1}$. This is obviously equivalent to maximizing expected future wealth for an indeterminate final period $T \gg t$. Finally, note that the hedge fund’s optimization problem does not consider any transaction costs (other than the per-position monitoring costs). This choice is only for parsimony and we don’t expect transaction costs to affect any of the predictions derived below.

2) Solution for a Specific Example

We now impose a specific structure on the nature of stock mispricings that will allow for an explicit, simple solution. We are convinced that the resulting predictions also obtain under a more general structure.

Specifically, we now assume that stock mispricings have the same magnitude and last for two periods. That is, in the period of occurrence, the mispricing is given by Δ ($-\Delta$). In the following period, the mispricing reduces to $\Delta(1 - \delta)$ with $0 < \delta < 1$ being the rate of alpha decay. After two periods, the mispricing is assumed to have disappeared. We further assume that every period t , a random number M_t of new stocks becomes mispriced (some positive, some negative). The number of mispriced stocks is small relative to the total number of stocks, $M_t \ll N$. It follows that in any period t , there are M_t “newly mispriced stocks” with an expected alpha of Δ , M_{t-1} “previously mispriced stocks” with an alpha of $\Delta(1 - \delta)$, and M_{t-2} stocks that just stopped being mispriced. Finally, we assume (without loss of generality) that $\bar{\sigma}^2 = \kappa_t \sigma_t^2$ for some $\kappa_t > 0$.

Recall that the hedge fund maximizes the expected return by choosing (i) in how many mispriced stocks to invest in and (ii) how much to invest in those stocks subject to not exceeding the volatility limit. For these choices, the hedge fund trades off the diversification benefits of investing into many mispriced stocks with the costs of monitoring a large number of open positions.

Let N_{At} be the hedge fund's choice of how many of the M_t newly mispriced stocks to invest in. Since newly mispriced stocks have the same maximum (absolute) mispricing, same volatility, and are uncorrelated with each other, the hedge fund will want to invest with equal (absolute) weights w_{At} into these N_{At} stocks.¹⁶ Similarly, let N_{Bt} be the number of previously mispriced stocks that the hedge fund chooses to invest in. Because they have the same level of (partially decayed) mispricing, the hedge fund will again want to invest with equal weights w_{Bt} into these N_{Bt} stocks.

Given these definitions, $N_{Pt} = N_{At} + N_{Bt}$ and the hedge fund's optimization problem can be written as

$$\max_{w_{At}, w_{Bt}, N_{At}, N_{Bt}} W_t [N_{At} w_{At} + N_{Bt} w_{Bt} (1 - \delta)] \Delta - c(N_{At} + N_{Bt}) \quad \text{s.t.} \quad N_{At} w_{At}^2 + N_{Bt} w_{Bt}^2 = \kappa.$$

It is easy to see that it will be suboptimal for the fund to choose $N_{At} < M_t$ while having $N_{Bt} > 0$. This is because all mispriced stocks have the same risk and the same monitoring costs, but newly mispriced stocks offer strictly higher returns. As such, the hedge fund will always want to prioritize investments into newly mispriced stocks (i.e., $N_{Bt} > 0$ only if $N_{At} = M_t$).

Proposition (Optimal Trading Rule):

Let $\bar{M}_{1t} \equiv \kappa_t \left(\frac{W_t \Delta}{2c} \right)^2 (1 - \delta)^4 - M_{t-1} (1 - \delta)^2$, $\bar{M}_{2t} \equiv \kappa_t \left(\frac{W_t \Delta}{2c} \right)^2 (1 - \delta)^4$, and $\bar{M}_{3t} \equiv \kappa_t \left(\frac{W_t \Delta}{2c} \right)^2$. The hedge fund's optimization problem has a unique solution which takes the following form:

- For $M_t \geq \bar{M}_{3t}$, the hedge fund only invests into some of the newly mispriced stocks. We have

$$w_{At} = \frac{2c}{W_t \Delta}, w_{Bt} = 0, N_{At} = \kappa_t \left(\frac{W_t \Delta}{2c} \right)^2, \text{ and } N_{Bt} = 0.$$

- For $\bar{M}_{2t} \leq M_t < \bar{M}_{3t}$, the hedge fund only invests into all newly mispriced stocks. We have

$$w_{At} = \sqrt{\frac{\kappa_t}{M_t}}, w_{Bt} = 0, N_{At} = M_t, \text{ and } N_{Bt} = 0.$$

- For $\bar{M}_{1t} \leq M_t < \bar{M}_{2t}$, the hedge fund invests into all newly mispriced stocks and some of the previously mispriced stocks. We have

$$w_{At} = \frac{2c}{W_t \Delta (1 - \delta)^2}, w_{Bt} = \frac{2c}{W_t \Delta (1 - \delta)}, N_{At} = M_t, \text{ and } N_{Bt} = \kappa_t \left(\frac{W_t \Delta (1 - \delta)}{2c} \right)^2 - \frac{M_t}{(1 - \delta)^2}.$$

- For $M_t < \bar{M}_{1t}$, the hedge fund invests into all newly mispriced stocks and all previously mispriced stocks. We have

$$w_{At} = \sqrt{\frac{\kappa_t}{M_t + M_{t-1} (1 - \delta)^2}}, w_{Bt} = \sqrt{\frac{\kappa_t (1 - \delta)^2}{M_t + M_{t-1} (1 - \delta)^2}}, N_{At} = M_t, \text{ and } N_{Bt} = M_{t-1}.$$

¹⁶ To see this, note that having equal weights as opposed to any others yields the same expected return, but minimizes the total variance of these investments.

Proof: We start with assuming that the hedge fund only invests into newly mispriced stocks. That is $N_{At} \leq M_t$, and $N_{Bt} = w_{Bt} = 0$. In this case, the Lagrangian of the fund's optimization problem becomes

$$\mathcal{L}(w_{At}, N_{At}, \lambda) \equiv W_t N_{At} w_{At} \Delta - c N_{At} - \lambda (N_{At} w_{At}^2 - \kappa_t),$$

where the Lagrange-multiplier λ needs to be positive. Solving the system of equations resulting from the first-order-conditions yields the unique solution

$$w_{At} = \frac{2c}{W_t \Delta}, N_{At} = \kappa_t \left(\frac{W_t \Delta}{2c} \right)^2, \text{ and } \lambda = \frac{W_t \Delta^2}{4c} > 0.$$

By assumption, $N_{At} \leq M_t$ and so this solution is only valid for $M_t \geq \kappa_t \left(\frac{W_t \Delta}{2c} \right)^2 \equiv \bar{M}_{3t}$.

Next, we consider the case where the hedge fund invests into all newly mispriced stocks, $N_{At} = M_t$, and chooses how many previously mispriced stocks to invest in. The Lagrangian is

$$\mathcal{L}(w_{At}, w_{Bt}, N_{Bt}, \lambda) \equiv W_t [M_t w_{At} + N_{Bt} w_{Bt} (1 - \delta)] \Delta - c (M_t + N_{Bt}) - \lambda (M_t w_{At}^2 + N_{Bt} w_{Bt}^2 - \kappa_t).$$

Solving the system of first-order-conditions again yields a unique solution with $\lambda > 0$:

$$w_{At} = \frac{2c}{W_t \Delta (1 - \delta)^2}, w_{Bt} = \frac{2c}{W_t \Delta (1 - \delta)}, N_{At} = M_t, N_{Bt} = \kappa_t \left(\frac{W_t \Delta (1 - \delta)}{2c} \right)^2 - \frac{M_t}{(1 - \delta)^2},$$

$$\text{and } \lambda = \frac{W_t \Delta^2 (1 - \delta)^2}{4c} > 0.$$

Clearly, we must have $0 < N_{Bt} \leq M_{t-1}$. These conditions imply that the solution is only valid in the range $\bar{M}_{1t} \equiv \kappa_t \left(\frac{W_t \Delta}{2c} \right)^2 (1 - \delta)^4 - M_{t-1} (1 - \delta)^2 \leq M_t < \kappa_t \left(\frac{W_t \Delta}{2c} \right)^2 (1 - \delta)^4 \equiv \bar{M}_{2t}$.

Since $\bar{M}_{2t} < \bar{M}_{3t}$, there is a range for M_t in which neither solution applies. This means that, for $\bar{M}_{2t} \leq M_t < \bar{M}_{3t}$, there exists neither an interior solution for N_{At} nor for N_{Bt} . We are thus left with a corner solution in which the hedge fund invests only in newly mispriced stocks but not in previously mispriced ones, $N_{At} = M_t$ and $N_{Bt} = 0$. w_{At} is then chosen to max out the volatility limit, yielding $w_{At} = \sqrt{\frac{\kappa_t}{M_t}}$.

Similarly, for $M_t < \bar{M}_{1t}$, there is another corner solution in which the hedge fund invests into all newly and previously mispriced stocks, $N_{At} = M_t$ and $N_{Bt} = M_{t-1}$. Finding the optimal w_{At} and w_{Bt} involves solving the first-order-conditions implied by the following Lagrangian:

$$\mathcal{L}(w_{At}, w_{Bt}, \lambda) \equiv W_t [M_t w_{At} + M_{t-1} w_{Bt} (1 - \delta)] \Delta - c (M_t + M_{t-1}) - \lambda (M_t w_{At}^2 + M_{t-1} w_{Bt}^2 - \kappa_t).$$

The unique solution is given by

$$w_{At} = \sqrt{\frac{\kappa_t}{M_t + M_{t-1} (1 - \delta)^2}}, w_{Bt} = \sqrt{\frac{\kappa_t (1 - \delta)^2}{M_t + M_{t-1} (1 - \delta)^2}}, \text{ and } \lambda = \frac{\Delta}{2} \sqrt{\frac{M_t + M_{t-1} (1 - \delta)^2}{\kappa_t}} > 0. \blacksquare$$

The optimal trading rule has intuitive properties. When M_t is very large, the hedge fund only invests in some of the newly mispriced stocks. The exact number of newly mispriced stocks into which it invests is increasing in the fund's wealth W_t , the level of the mispricing Δ , the volatility limit κ_t , and decreasing in the monitoring cost c . For a lower M_t , there is first a range in which the hedge fund only invests in all newly mispriced stocks, but not in previously mispriced ones. As M_t gets lower still, the hedge fund also starts investing into previously mispriced stocks, where the number of such positions is again an increasing function of wealth W_t , mispricing Δ , volatility limit κ_t , and decreasing in monitoring cost c . Finally, when M_t and M_{t-1} are very low, the hedge fund invests into all newly and previously mispriced stocks.

3) Life-cycle of a Round-trip Trade

We now describe the life-cycle of a round-trip trade—i.e., its *initiation*, the *rebalancing* and its *closure*. Consider a new mispricing in stock i occurring in period t . In that period, the hedge fund *initiates* the trade by investing $W_t \times w_{At}$ of risk capital into that stock.¹⁷ Depending on whether this is an under- or overpricing, this would take the form of either a long or a short position. The hedge fund *closes* its position after either one or two periods, depending on how many newly mispriced stocks there will be in the next period ($t+1$). All intermediate trades are defined as *rebalancing* trades.¹⁸ (These rebalancing trades will typically result in a gradual downscaling of the position concomitant to the decay in alpha.)

4) Empirical Predictions

Let N_{At-1} be the number of positions in stocks that became mispriced in $t-1$ ($N_{At-1} \leq M_{t-1}$) and let N_{Bt-1} be the number of positions in stocks that became mispriced in $t-2$ ($N_{Bt-1} \leq M_{t-2}$).

In period t , the hedge fund closes all the N_{Bt-1} positions in stocks that become mispriced in $t-2$ (because they stop being mispriced). In addition, the hedge fund may need to close some (and perhaps all) of its N_{At-1} positions in stocks that become mispriced in $t-1$. Specifically, out of the N_{At-1} positions, it will only be able to hold on to N_{Bt} positions, where typically $N_{Bt} < N_{At-1}$.¹⁹ Each of these prematurely closed positions is followed by a positive expected return ($\Delta(1 - \delta)$). As such, the average return after closing trades in period t is given by $X_t \Delta(1 - \delta)$, where X_t is defined as

$$X_t \equiv \frac{(\text{Max}\{N_{At-1} - N_{Bt}, 0\})w_{At-1}}{(\text{Max}\{N_{At-1} - N_{Bt}, 0\})w_{At-1} + N_{Bt-1}w_{Bt-1}}.$$

Note that X_t is a fraction; i.e., $0 \leq X_t \leq 1$. The fraction becomes zero if $N_{At-1} \leq N_{Bt}$, which should be the exception rather than the rule (see footnote 6). Otherwise it will be strictly positive. Moreover, as long as N_{Bt-1} is not zero, fraction X_t is strictly less than one, in which case it is decreasing in N_{Bt} .

¹⁷ When there are too many newly mispriced stocks, the hedge fund cannot invest into all of them. When this case occurs, we simply assume that the stock under consideration is one of the N_{At} newly mispriced stocks in which the hedge fund does invest.

¹⁸ Note that, with this definition, only positions that are open for two periods can have a rebalancing trade (in the intermediate period).

¹⁹ To see why typically $N_{Bt} < N_{At-1}$, note that the hedge fund always prioritizes newly mispriced stocks (because of the alpha decay in previously mispriced stocks). Thus, as long as the number of open positions $N_{Pt} = N_{At} + N_{Bt}$ does not drastically increase from one period to the next, the hedge fund ends up closing some positions in previously mispriced stocks to shift the risk capital into newly mispriced stocks.

Finally, note that all terms entering X_t except for N_{Bt} are pre-determined (i.e., depend on parameters from period $t-1$). Thus, only N_{Bt} matters for the description of the relationship between contemporaneous characteristics (such as M_t or W_t) and post-closure returns.

The following empirical predictions follow immediately:

Prediction 1: *The initiation of a trade is more predictive of future returns than the closing of a trade.*

Proof: Initiating trades in any period t are followed by an average return close to Δ .²⁰ Closing trades are followed by an average return of $X_t\Delta(1 - \delta)$, where $0 \leq X_t \leq 1$ (see above). ■

Prediction 2: *The closing of a trade is followed by future returns in the opposite direction of the closing trade. In other words, the difference in post-closure returns between closed long and short positions is positive—implying that the hedge fund “leaves money on the table”.*

Proof: For underpriced stocks ($\Delta_{it} = \Delta$), the hedge fund took long positions, which require selling at closure. Yet, as seen above, the average return following such long sells is positive. For overpriced stocks ($\Delta_{it} = -\Delta$), the hedge fund took short positions, which require buying at closure. The average return following such short buys is negative. ■

Prediction 3: *The return difference between closed long and short positions should be higher in periods when lots of new stock mispricings occur (and thus when lots of new trades are initiated).*

Proof: X_t is decreasing in N_{Bt} , which in turn is decreasing in M_t . ■

Prediction 4: *The return difference between closed long and short positions should be higher after periods in which the hedge fund has had low returns.*

Proof: X_t is decreasing in N_{Bt} , which in turn is increasing in r_{Pt} (through W_t). ■

Prediction 5: *The return difference between closed long and short positions should be higher in periods when stocks are more volatile.*

Proof: X_t is decreasing in N_{Bt} , which is increasing in κ_t and thus decreasing in σ_t . ■

The intuitions for these predictions are straightforward. Prediction 1 follows from the fact that the hedge fund initiates positions when the mispricing has just occurred and is thus the biggest, whereas it closes its positions when the mispricing has (partially or fully) decayed. Prediction 2 says that hedge fund “leave money on the table”; i.e., it could have made additional profits from holding on to its positions for longer. This result naturally follows from the fund’s desire to limit total position monitoring costs, as it may induce the fund to close positions in partially mispriced stocks when better investment opportunities become available. Finally, predictions 3 to 5 say that such position closures of still partially mispriced stocks occur more often when there are more new mispricings, when the fund has suffered from poor returns, or when the volatility constraint becomes more binding due to an increase in stock return volatility.

²⁰ Occasionally, it can occur that a new position is opened in a previously mispriced stocks. Specifically, when the hedge fund has had a phenomenally high return, it chooses to open many new positions, which may entail opening a position in a previously mispriced stock that the fund had not yet invested in. Since such position openings will be rare, the average post-initiation return will be strictly larger than $\Delta(1 - \delta)$.