The Pennsylvania State University

The Graduate School

Smeal College of Business

ESSAYS ON THE STRUCTURE, PERFORMANCE, AND BEHAVIOR OF HEDGE FUNDS

A Dissertation in

Business Administration

by

Grant V. Farnsworth

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Doctor of Philosophy

December 2015

The dissertation of Grant V. Farnsworth was reviewed and approved* by the following:

Timothy Simin
Associate Professor of Finance
Smeal Research Fellow in Finance
Dissertation Adviser
Chair of Committee

Jared Williams
Assistant Professor of Finance

Matthew Gustafson
Assistant Professor of Finance

Quanwei Charles Cao
Smeal Chair Professor of Finance

Jeremiah Green
Assistant Professor of Accounting

Brent Ambrose
Smeal Professor of Risk Management
Director, Institute for Real Estate Studies
Director, Ph.D. Program

* Signatures are on file in the Graduate School

Abstract

This dissertation contains three essays on the structure of hedge funds and the behavior of hedge fund managers and investors. In the first, I study the dynamics of changes in the leverage choices made by hedge fund managers and the implications for investors. Using SEC and commercial hedge fund database information, I examine the time-series and cross section of hedge fund leverage and find that, overall, hedge fund leverage changes are not driven by certain return-driven motivations, such as levering up before profitable times or in order to take advantage of concentrated investment opportunities. Instead, hedge fund leverage changes are driven primarily by risk mitigation measures, which may be imposed by external actors, such as prime brokers. Hedge funds lever down during volatile times and when risk spreads are high. Interestingly, when hedge fund leverage is perturbed by fund flows, I find no evidence that managers trade to return to a target leverage within four quarters. Thus overall hedge fund leverage may be the result of the fund's flow history, rather than predetermined leverage targets. In the second study, I examine hedge fund inceptions. New hedge funds may come into being because managers observe high demand for certain types of funds and create them to absorb this demand, or alternatively as a result of an innovative investment idea. I create empirical proxies that allow me to distinguish the two types of inceptions and show that funds that came about because of the supply of managerial investment from those that came about because of investor demand. Inceptions of the former type outperform those of the latter by a significant 4 to 5% annually over the first five years. In my third study, I use a special environment, a platform of hedge fund separate accounts, as a laboratory to examine the role of share restrictions and third-party evaluation in hedge fund returns. Separate accounts are tied to existing funds but have much lower share restrictions and feature third-party return evaluation. By comparing separate account and main fund returns for the same funds, I find that a reduction in share restrictions leads to a performance penalty of 1.7% per year. Also, the high liquidity and third party evaluation leads to reported returns that feature 33% less serial correlation. This latter result suggests that managers in the main fund use discretion in reporting practices to induce serial correlation in the reported returns of their main fund, which would lead to artificially good performance in risk-adjusted evaluation metrics for those funds.

TABLE OF CONTENTS

List of Figures

List of Tables

Preface

The second chapter is based on work that is coauthored with Charles Cao and Hong Zhang. The third chapter is based on work that is coauthored with Charles Cao, Bing Liang, and Andrew Lo.

Chapter 1: Strategic Hedge Fund Leverage and Investor Welfare: A Holdings-Based Approach

Abstract

Dynamic use of leverage is a distinguishing characteristic of hedge funds, but does this leverage benefit investors? In this paper I test whether hedge funds' use of leverage successfully improves risk-adjusted performance. I use a novel, holdings-based, approach to construct a panel of leverage changes and find that hedge fund leverage changes are not driven by return-centric strategies such as levering up during or before profitable times or to pursue exceptional opportunities. However, I do find that leverage decisions are based in part on risk targeting concerns. Surprisingly, when leverage is mechanically perturbed by fund flows, managers do not trade back toward original levels, indicating that leverage is driven more by funds' history than by managers' skill at choosing a leverage target.

Keywords: hedge funds, leverage, timing, holdings, borrowing constraints
JEL classifications: G01, G23, G11

1.1 Introduction

Hedge funds are unique among financial intermediaries in that their primary objective is the employ of sophisticated trading strategies and superior investing skills. Inasmuch as hedge funds do, indeed, generate excess risk-adjusted returns for their investors, what is the source of these returns? A defining characteristic and primary advantage of hedge funds is that they are free to utilize time-varying leverage with little oversight and few constraints. They may, therefore, use this leverage to the benefit of their investors, for example by employing leverage as a market-timing mechanism or to take full advantage of profitable opportunities as they arise. At the same time, the nonlinear fee structure, which gives hedge fund managers a large fraction of upside returns without imposing downside costs, may incentivize perverse leverage choices that benefit managers at the expense of investors.

In the absence of time-series information about hedge fund leverage, investors and regulators do not know whether they would approve of the manner in which hedge fund managers change leverage over time. In fact, it has been suggested, by Lo (2008) among others, that increasing hedge fund leverage over time in pursuit of high returns poses a systemic threat to the stability of the entire financial system. Certainly, given the lack of disclosure on the issue, investors can have little clue about when hedge funds change their leverage. In the absence of this knowledge, investors cannot rebalance their overall portfolios in response to changes in the effective risk exposure of their hedge fund holdings. In short, the secretive nature of hedge funds means that the dynamics of hedge fund leverage are almost completely unavailable to researchers, regulators, and investors, despite their importance to policy and investment decisions.

In this paper I present a new method for constructing the full panel of hedge fund leverage using a combination of information reported to commercial hedge fund databases and holdings information required by the SEC in hedge funds' quarterly 13(f) filings. I use the resulting leverage estimates to test whether hedge funds effectively use their leverage either to market-time their investments or to take full advantage of profitable investment opportunities. I further test for alternative drivers of hedge fund leverage changes, such as a desire to maintain raw returns in the face of increasing competition. In addition, I examine whether hedge funds pursue leverage targets by testing whether they rebalance back toward their original level when leverage is mechanically perturbed by fund flows.

Mechanically, hedge funds typically utilize the services of prime brokers—often investment banks—who serve as custodians and provide operational support and financing. These prime brokerages permit several forms of hedge fund leverage. In addition to individually negotiated loans, prime brokers typically provide a flexible line of credit to client hedge funds, which permits rapid changes in effective leverage without prior notice. Since hedge funds are also free to engage in short-selling, funds can also sell assets short and use the generated funds to increase profitable long positions, which may have very different risk exposures than their shorts. While this practice may differ from the most intuitive concept of leverage, it is nonetheless commonplace in the industry and frequently referred to as "leverage" by managers and prime brokers. Because my measure of hedge fund leverage incorporates the total value of long assets held by the fund, it accounts for both types of leverage.

Prime brokers also keep a watchful eye on the risk exposure of client hedge fund portfolios. They decide the explicit leverage limits the funds face as well as the cost of leverage. A typical prime broker has a risk-monitoring committee that meets regularly to discuss whether client hedge funds are maintaining acceptable risk in their portfolios. Concerns about a fund's portfolio can lead to requests to reduce risk, increased borrowing costs, or hard limits on borrowing. Hedge funds often have internal risk controls as well, which are put in place to give brokers and investors a measure of confidence.

In addition to the concerns of investors and counterparties, changes in hedge fund leverage are of particular interest to regulators and participants seeking stability of financial markets. Historically, several well-known funds have failed at least in part because they employed a large amount of implicit or explicit leverage; Long Term Capital Management (LTCM) in 1998 and two of Bear Sterns' credit strategy funds in 2007 are examples. LTCM, in particular, demonstrated that even individual fund failures can have wide-ranging systemic effects on the economy. The largest of today's hedge funds are much larger than LTCM was at its largest[1]. Additionally, the number of hedge funds in operation has increased dramatically over the last two decades. Given these facts, a rise in hedge fund leverage could increase systemic risk in the economy dramatically without regulators and counterparties observing the change. An understanding of what drives leverage changes is therefore timely and valuable even for individuals wholly outside of the hedge fund industry.

Hedge fund managers take as an incentive fee a large and fixed proportion (typically 20%) of the total *increase* in the value of the fund. This gives the manager's payoff the appearance of a call option: in the presence of positive returns hedge fund managers' compensation rises to lavish levels while the downside costs are limited. Since leverage increases the risk exposure and variance of fund returns, it also increases the optional value of the manager's fees. On the other hand, high leverage can also bring dramatic losses to the fund and induce investors to withdraw their capital. The manager must therefore balance the expected benefits of leverage against its costs. Some investors may actually prefer some degree of leverage in a hedge fund even though it more than proportionally increases the manager's fees. For example, an investor unable to obtain cheap credit or unable to take explicitly levered positions (certain institutional investors, for example) might prefer a hedge fund in their portfolio to use leverage. Clientele and asset effects, therefore, may generate a static pattern in the cross-section of hedge fund leverage.

In the time-series, changes in leverage are the primary variable of interest. What pattern of leverage changes would an investor, then, desire? Certainly if a hedge fund manager is able to correctly predict that her current portfolio will have positive returns with high certainty, an investor would want the manager to magnify these returns through leverage. Similarly, when uncertainty about the fund's assets is higher, an investor would be likely to prefer less risk in the hedge fund portfolio. Additionally, some funds specialize in taking

[1]LTCM's assets under management were about $4.74 billion in the beginning of 1998. As of this writing, Bridgewater Associates has assets under management of about $122 billion.

advantage of rare opportunities, such as undervalued assets and arbitrage opportunities. For funds of this type, a desirable pattern would be to maintain moderate or no leverage during normal periods and then to lever up when a large opportunity presents itself. This pattern of leverage would tend to increase the fund's portfolio concentration (reduce diversification) during times of increasing leverage as the borrowed capital is invested primarily in the new opportunity.

Investors may be less happy with other patterns of leverage changes. For example, managers have incentive to maintain numerically large returns in order to maximize their fees and attract the attention of new investors. However, as hedge fund profitability often depends on arbitrage opportunities of fixed size, hedge funds face diseconomies of scale and the entry of competing funds can also drive returns down. A manager desiring numerically high raw returns may be tempted to increase leverage, magnifying the fund's risk contrary to the interests of the fund's investors. Since there has been tremendous entry in the hedge fund industry overall, this would imply increasing leverage across the industry.

Another leverage pattern that might not please investors and regulators would be for the fund manager to fail to actively monitor and maintain leverage. Many factors affect a fund's leverage over time without the manager's explicit intervention. For example, if the manager does not trade, fund leverage changes mechanically with the performance of existing assets (positive performance increases the assets under management without increasing the debt-like obligations) or with investor flows. If the manager does not actively monitor leverage and trade to a desired leverage level, these factors can potentially move hedge fund leverage around in a pattern completely unrelated to the manager's confidence in the portfolio. In fact, given the positive serial correlation in hedge fund returns, both flows and positive performance mechanically reduce leverage precisely during times when per-dollar returns for the fund are likely to be high. If the manager fails to actively maintain desired leverage targets, the natural leverage pattern works against investor interests.

I find that, overall, increases in hedge fund leverage are not associated with subsequent positive risk-adjusted performance, implying that managers do not successfully use leverage to time the returns of their portfolio assets. However, hedge funds do decrease leverage during times of high market volatility (both forward and backward looking), consistent with risk targeting behavior. Hedge funds do not tend to increase leverage following poor performance or increased competition, nor has there been a trend toward increased leverage over time. Further, hedge funds tend to increase their diversification when levering up, so funds do not appear to lever up to take advantage of rare and valuable investment opportunities. I find that hedge funds lever down when credit spreads are high, either because of cost considerations or because their prime brokers request reduced leverage during these times. Surprisingly, when hedge fund leverage is mechanically perturbed by fund flows, hedge fund managers do not rebalance back to the original levels, implying that hedge fund leverage may be largely a result of the history of flows and performance of the fund, rather than a conscious choice by fund management.

The rest of the paper is organized as follows. Section 1.2 summarizes related literature. Section 1.3 discusses possible causes of changes in hedge fund leverage and develops testable hypotheses. Section 1.4 describes the data and methodology used to construct the leverage panel. Section 1.5 discusses the results of the analysis. Section 1.6 concludes.

1.2 Related Literature

The need for research on hedge fund leverage has not been lost on the field. However, it has been difficult to circumvent the difficulties of hedge fund data. In the past, researchers have primarily used estimates (Lo, 2008) or employed indirect techniques. Primary among indirect techniques has been regression analysis on survey data (Eichengreen and Park, 2002) or of returns on macroeconomic factors (McGuire and Tsatsaronis, 2008). The latter find the unexpected result that the estimated leverage for such diverse strategies as fund of funds, event driven funds, and equity hedge funds move largely together, especially their exposure to the broad market and SMB factors. This suggests that a common cause, such as credit availability or perceived riskiness of the market, is at work.

Some commercial hedge fund databases request that the funds specify some attributes of their leverage usage, for example the gross "leverage" field in CISDM or the "average leverage" and "maximum leverage" fields in Lipper TASS. The latter is used by Schneeweis, Martin, Kazemi, and Kawavas (2005), who examine risk-adjusted returns for levered and less levered funds but find little difference. While these fields can provide valuable information on a fund-by-fund basis, they are static snapshots of leverage and also self-reported. The leverage field for TASS, for example, contains a variety of answers that seem to indicate that the question is not being interpreted the same by all managers. Even if all managers did honestly report their leverage in a uniform manner, the time-varying nature of leverage limits the value of these reported values. Liang and Qiu (2013) analyze changes in this data field by comparing reported leverage from many sequential versions of TASS purchased over 9 years. Unfortunately they find that few funds modify these fields over time.

One of the most significant contributions to our understanding of hedge fund leverage comes from Ang, Gorovyy, and van Inwegen (2011). The authors obtain access to leverage information from a fund of hedge funds and are able to examine a panel of leverage for that set of funds. In their sample, the average long-only leverage over all hedge funds is 1.36 and long-only leverage is almost perfectly correlated with gross leverage. They also find significant persistence in leverage and that hedge fund leverage is generally countercyclical with respect to the market leverage of listed financial intermediaries. Broadly speaking, they do not identify fund-specific factors that significantly predict changes to leverage (aside from past leverage).

The use of SEC 13(f) filings to examine the holdings of financial institutions has been increasingly popular, though to my knowledge it has not been previously used to examine leverage. For example, Brunnermeier and Nagel (2004) use holdings information to show that hedge funds in the late 1990s tended to ride the technology bubble, rather than exert significant correcting price pressure by attacking it.

In the leverage literature, Fostel and Geanokoplos (2008) predict that leverage in financial institutions should decrease during times of high volatility and uncertainty. Acharya and Viswanathan (2011) show that adverse market conditions induce sharply decreasing leverage and sudden drying up of market liquidity. Adrian and Shin (2010) examine the time series of margin requirements (haircuts), treating it as a measure of leverage, and show that leverage is strongly procyclical. Stein (2009) explores how hedge funds choose optimal leverage from their perspective and create a fire-sale externality (raising the likelihood of a severe crash). Fund managers may be less concerned about the economy-wide implications of their leverage choices than the implications for their own fund. Duffie, Wang, and Wang (2008) describe a theoretical framework to show that fund managers trade off the benefits of higher leverage against the costs of adjusting that leverage during bad times to derive an optimal leverage strategy that depends on the transactions cost of changing that leverage. Dai and Sundaresan (2010) show that fund managers have a short funding option both on the investor side and the borrowing side. Managers then choose an optimal leverage ratio that balances the expected gains against the possibility of investor withdrawals or the imposition of funding limits from their lenders.

1.3 Hypothesis Development

1.3.1 Market Timing

Hedge fund managers rely on skill to identify opportunities for investments that generate excess risk-adjusted returns for their investors. Typically this skill involves execution of trading strategies and/or investment in undervalued assets. However, hedge fund managers also have leverage as one of their investment tools. If a manager can predict that the fund's portfolio has a high probability of producing positive returns, an appropriate strategy would be to increase leverage. If hedge fund managers do, in fact, have the ability to estimate a time-series of confidence in their investments, then the resulting time-series of leverage should successfully time the portfolio. This suggests the following hypothesis:

Hypothesis 1.1 *If hedge funds successfully use leverage to take advantage of the ability to* *predict positive fund returns, then there should be a positive relation between risk-adjusted returns and contemporaneous or subsequent leverage changes.*

As an alternative to this hypothesis, if managers cannot time their fund returns but do monitor leverage to maintain targets, then one would not expect a relation between leverage and returns. Alternatively, if managers do not trade to maintain leverage, then there will be a negative relation between contemporaneous leverage changes and returns because returns increase investor capital, reducing contemporaneous leverage. This last case does not predict a relation between leverage and subsequent returns.

1.3.2 Portfolio Concentration

Many hedge fund trading strategies are opportunistic by nature. Arbitrage opportunities and mispriced assets may not be frequent events for some hedge funds. If this is the case, a reasonable leverage strategy would be to maintain little or no leverage during normal times, but to lever up as much as possible to take advantage when profitable opportunities arise. If hedge funds use leverage to load up on single-asset opportunities, whether they are successful in earning profits on those opportunities or not, we would expect portfolio concentration to increase when leverage increases. This leads to the following hypothesis:

Hypothesis 1.2 *If hedge funds lever up to take maximum advantage of large opportunities, leverage changes will be positively related to contemporaneous changes in portfolio concentration.*

Alternatively, if hedge funds modify leverage for reasons unrelated to the quality of investment opportunities, we might expect concentration to remain the same or even fall as leverage increases because economies of scale make diversification more cost effective when the fund's position is large.

1.3.3 Volatility/Risk Targeting

A hedge fund manager must choose leverage to balance the benefits of greater exposure against the costs. Dai and Sundaresan (2010) describe how increasing leverage can lead to capital withdrawals by investors, a negative outcome for the manager. To prevent this, the manager may list specific volatility targets in the explicit objectives of the fund in the fund's investor documents. This would increase investor confidence in the risk controls of the fund and prevent capital flight. Alternatively, hedge fund managers may have implicit volatility or risk targets for the fund that target what they feel is an upper-bound of the acceptable risk. As an example, a manager using mean-variance optimization as part of their allocation scheme would likely have a fixed volatility target that would enforce lower leverage during times when asset volatility and uncertainty is high. Alternatively, a fund may have less formal internal risk monitoring systems that suggest reduction in exposure when asset risk is high. This line of thinking suggests the following hypothesis:

Hypothesis 1.3 *The change in leverage for a hedge fund management company will depend negatively on changes in the volatility and risk environment.*

1.3.4 Credit Constraints and Costs

One of the more common explanations for changes in aggregate hedge fund leverage is that hedge funds lever up and down in response to the availability of credit. Credit availability can be proxied by spreads in interest rates. If a hedge fund's leverage is being moderated by risk controls, whether they be internal controls or whether the prime broker is imposing precautionary leverage limitations or costs, we would see hedge fund leverage fall when credit spreads rise.

Alternatively, hedge funds could be constrained in the amount of credit they can *affordably* take. Positive leverage increases the riskiness of a hedge fund, so incremental explicit costs of borrowing may be increasing as well. Since the availability of capital for lending by prime brokers and other lending institutions varies over time with the economy-wide financial situation, the optimal amount of leverage might be expected to vary with those conditions. This is essentially the assumption made by Adrian and Shin (2010), when they treat the time-series of margin constraints as a proxy for the time-series of leverage taken by financial institutions. Indeed, this leads to the following hypothesis:

Hypothesis 1.4 *If lending costs and constraints bind, changes in hedge fund leverage should depend negatively on measures of the cost and difficulty of borrowing.*

Explicit costs of borrowing can be measured in terms of interest rates. I decompose interest rates for a risky borrower into the time-value of money (the risk-free rate) and the spread imposed by the lender as compensation for default risk. More specifically, I examine changes in the risk-free rate (proxied by one-month libor), the TED spread (three month libor minus three month treasury yields), and the BAA-AAA credit spread, published by the Federal Reserve Board. Although it may not be clear how risky lenders consider typical hedge funds to be, given the high turnover of hedge funds, one can be confident that they have significant credit risk. Thus, changes in the cost of credit in the bond market are a proxy for changes to the cost and availability of credit to hedge funds. I ignore the term spread here because hedge fund borrowing is typically in the form of short-term contracts or lines of credit based on libor plus a spread.

1.3.5 The Effect of Flows

Leverage is defined by the ratio of the value of portfolio assets to the investor capital of the fund. This means that if a manager does not change the investments of a fund over time and the fund experiences positive or negative investor flows, the leverage ratio of the fund will change mechanically. Positive investor flows will reduce fund leverage and negative flows will increase it. One might expect that a manager's target leverage ratio and available credit would scale with the size of the investor capital of the fund. However, many hedge fund positions do not lend themselves to rapid investment and liquidation. A hedge fund holding distressed assets, for example, may not find it prudent to sell off these assets until their prices stabilize or until suitable liquidity is available. Similarly, some hedge funds hold large indivisible assets, such as real estate property or whole enterprises. Whether it is costly or impossible to change positions, hedge funds in this situation will find that investor flows modify leverage levels. Alternatively, if hedge fund managers do not actively follow leverage targets, perturbations in leverage due to investor flows may cause permanent changes in fund leverage. These ideas can be tested with the following hypothesis:

Hypothesis 1.5 *If hedge fund managers are able to actively pursue leverage targets, then when flows perturb leverage, fund managers should trade to return leverage to its former level.*

This hypothesis has a different implication for different horizons. If the manager can rebalance back to target levels within a quarter, there should be no relation between flows and contemporaneous or subsequent leverage changes. Alternatively, if the manager takes months to return to target levels, there would be a negative contemporaneous relation between flows and leverage and a positive relation between flows and subsequent leverage changes.

1.3.6 Profit and Competition

The dramatic increase in the number and size of hedge funds over time has led researchers to rationally question whether these funds are beginning to crowd each other out as they seek to trade on the same set of opportunities. In the presence of increased competition, the per-trade or per-dollar return for funds is likely to decrease. If that is the case, how might managers react? Remembering that hedge fund managers have nonlinear payout structures and therefore benefit disproportionately from numerically high returns (even if volatility is high as well), one might expect managers to lever up so that the amount of invested money is greater and the returns per investor equity are maintained. This incentive would be especially strong for hedge funds that have hurdle rates such that the manager does not receive any incentive fee until a nonzero positive return hurdle has been surpassed.

Besides the conditions governing incentive fee calculation, many hedge fund managers are also interested in expanding the total assets under management of the fund. Since investors have many funds to choose from and in the absence of full transparency must rely on little more than historical returns to choose between hedge funds, managers have incentive to provide a stream of returns that is eye-catching and attractive to potential investors. A downward trend in signed returns, even if the returns are positive, will tend not to attract new money. Hedge fund managers in this position would have incentive to increase leverage to drive up returns, even if it means increasing the riskiness of the fund. This leads to the following hypothesis:

Hypothesis 1.6 *If hedge funds use leverage to maintain return targets in the face of competition, entry in a hedge fund strategy should lead to subsequent increases in leverage for funds within that strategy.*

There has been a significant increase in the number of hedge funds in operation over time. However, fund entry is higher in some periods than in others and investor interest changes over time by hedge fund investment strategy. Every fund that reports to one of the commercial hedge fund databases also states its inception date as well as its strategy, so it is possible to compute the number of new funds per strategy over time. If increased competition drives per-dollar-invested hedge fund returns down and maintaining return levels is an important determinant of leverage changes, then funds in strategies experiencing high entry may subsequently increase their leverage. I therefore add the following hypothesis:

Hypothesis 1.7 *If hedge funds use leverage to maintain return targets in the face of competitive pressure, low returns for a period of time should lead to increases in hedge fund leverage.*

If the factor driving the relation between competition and increased leverage is a fall in per-dollar returns, it may take some time for funds to react to falling profitability (it may take time for new funds to invest enough to make a difference in the market environment, and then it may also require time for the incumbent fund managers to observe their falling returns). In this case there would be a delayed negative relation between hedge fund returns and leverage. That is, a hedge fund generating protracted poor returns might choose to lever up and take more aggressive positions to maintain a target return. If this is the case, there will be a negative relation between past returns and changes in leverage.

1.4 Hedge Fund Data

This section describes the data sources and methodology used to construct the final leverage sample. I also discuss how leverage will be defined and used throughout the analysis.

1.4.1 Hedge Fund Holdings and Commercial Databases

Holdings information for all institutional investment managers is taken from Thomson Reuters Institutional 13(f) Holdings s34 master file. This file contains the holdings of all institutions that file form 13(f) holdings with the SEC. Since 1978, the SEC has required that any manager that uses U.S. mail in the course of its business (regardless of domicile and clientele) and has investment discretion over \$100 million or more of 13(f) securities must report its holdings in these securities quarterly. These securities include long positions in stock as well as long put and call options on stocks, though option information is frequently not found in the s34 file. It should be noted that other securities, such as futures contracts and short positions on equity, are not included in the list of 13(f) securities and may not be reported. As of each quarter end, I compute the market value of all reported securities for all institutional money managers. Because Thomson's 13(f) filings do not contain information about whether the institution is a hedge fund management company, I perform this step for all 13(f) companies.

I use three commercial hedge fund databases to obtain the time-series of assets under management (AUM) for each reporting hedge fund: Lipper TASS, HFR, and BarclayHedge. These databases provide information about the management company for each reporting fund, but there is no unique identifier to match them to the 13(f) companies or to companies appearing in more than one hedge fund database. For this reason I ultimately construct three separate mappings from the commercial databases to the 13(f) data. The AUM for each fund gives the total value of the investor capital in the fund on a monthly basis. Throughout this paper I will refer to the value of the investor capital as AUM and I will refer to the (possibly levered) value of the 13(f) assets held by a fund as the "asset value." The relation between these two quantities will ultimately provide a leverage measure.

Funds in the commercial databases report returns consistently but reported AUM numbers have several potential issues. The first is that funds sometimes report the exact same AUM several months in a row (perhaps because the manager doesn't feel like updating the number or because the reported AUM is rounded and the AUM hasn't deviated very far). To handle this, I delete reported AUM numbers that are identical to what was reported in the previous month. Checking empirically, this change has no measurable effect on any of this paper's results but I implement it as a best practice on principle. Funds sometimes also have gaps in their reported AUM series. Because this project involves monitoring changes in AUM and the value of held positions, either ignoring funds or months with missing AUM or using straight-line interpolation to estimate missing AUM would be inappropriate. Instead, when a gap in the AUM series appears I use returns and available AUM numbers to estimate dollar flows to the fund over the gap. I then reconstruct the missing AUM using monthly returns and spreading the multi-period dollar flows equally among the months for which AUM is missing. For example, if the AUM for a particular fund is missing in months $t+1$ through $t+k-1$ inclusive, then total flows over that period can be estimated using

$$\text{Flow}_{t,t+k} = AUM_{t+k} - AUM_t \cdot \prod_{i=1}^{k}(1+r_{t+i}). \tag{1}$$

This is derived from the approach taken by Chevalier and Ellison (1997), Sirri and Tufano (1998), and Agarwal, Daniel, and Naik (2004), and modified to give the dollar flow instead of percentage flow. I then estimate the missing AUM for each $j \in [1, k-1]$ as

$$\widehat{AUM}_{t+j} = AUM_t \cdot \prod_{i=1}^{j}(1+r_{t+i}) + \frac{j \cdot \text{Flow}_{t,t+k}}{k}. \tag{2}$$

I use this approach to fill gaps in the AUM series of up to 12 months. If the hole is at the beginning or end of the reported series, I use the expression in Equation (2), omitting the Flow term.

After compiling AUM and holdings value information for each management company within each database, I use a hybrid computer/human approach to construct three mappings, one from each commercial database to the set of 13(f) institutions associated with each hedge fund management company. To perform this nontrivial operation, I use the management company name, location, dates of operation, and AUM. Fuzzy matching using substring analysis, generalized Levenshtein distance, and other company characteristics are used to provide a list of 13(f) institution possibilities for each commercial database management company and positive matches are identified by hand with a bias toward excluding false positives. Because managers change the way they report their name over time, both in the hedge fund databases and 13(f) data, all possible names for each manager ID in both databases are compared against each other when finding candidates.

Commercial database-reported AUM and Thomson holdings asset value are then merged by institution and date for each match. Note that many academic studies that use 13(f) holdings to examine hedge funds must perform a separate step to identify which of the 13(f) institutions are primarily hedge fund management institutions. My approach yields this information as a byproduct of the commercial hedge fund database to 13(f) merge. I find all 13(f) institutions that are reported in the commercial databases as being management companies and then the filters I impose later naturally remove institutions that manage hedge funds only as a small part of their business, such as certain investment banks.

There are several implicit filters in the merged sample. For example, since the SEC requires the filing of a 13(f) only if the management company has $100M worth of US assets[2], small hedge funds who choose not to report to the SEC may not be represented in the 13(f) sample. Along similar lines, hedge fund reporting to the commercial databases is voluntary, though it is an attractive option, even for small funds. Appearing in commercial

[2]Virtually all US-listed equities are on the SEC list, as are many equity options.

databases may help to attract investor attention and establish their reputation.

I also use explicit filters to remove funds that are not good candidate matches or whose assets do not lend themselves to this paper's analysis. For example, some 13(f) institutions are associated with many funds or large non-hedge-fund activities. For these institutions, the 13(f) holdings may be much larger than the sum of the reporting funds managed by them. Also, some funds specialize in non-reportable assets, such as futures or other non-reportable derivative contracts. To exclude these types of matches, I require that the ratio of 13(f) reported asset values to summed reported hedge fund assets under management be no less than 1/5 and no greater than 5. This does exclude some valid matches that use extremely high or extremely low leverage, but since my analysis is primarily concerned with changes in leverage for typical funds, this exclusion seems justified.

There is significant overlap between the institutions matched to each of the commercial databases; many funds report to more than one database. Figure 1.2 illustrates the number of management companies from the 13(f) data were matched to each of the commercial databases. In order to reconcile these overlaps, for each 13(f) company I choose the commercial database with the longest history and fewest missing values.

For reference, there are 7,100 total managers in the BarclayHedge universe, 6,833 of which report a valid AUM. Of these, 2,543 report more than $100 million USD of assets under management at some point in their history (though with leverage, their position value may be higher). Of these larger managers, 1,501 are domiciled in the United states. Note that this does not mean only these managers are candidates for 13(f) reporting; the SEC requires all institutions with $100 million worth of 13(f) securities to report, regardless of their domicile.

After completing the merge and removing remaining missing values in the leverage stream, there are a total of 31,395 quarterly observations from 1,322 managers with an average of 3.75 funds per management company in the sample.

Because a management company may run several funds, the possibility of a management company having funds in several strategies is high. Further, there are many hedge fund strategies that do not lend themselves to my analysis because they have the possibility of investing heavily in assets that are not 13(f) reportable. The 13(f) filings contain only information about long equity and equity-like holdings. I therefore would like to restrict my primary analysis to managers that run funds in equity-oriented strategies. Although the commercial hedge fund databases use different strategy classifications, they each have a fairly large category designating equity long/short funds, which I choose for my restricted sample. Because managers may have more than one fund and therefore more than one strategy, I further restrict the sample to managers that run funds only of the long/short equity variety. It should be noted that though most management companies run more than a single fund, it is very common for all funds to report the same hedge fund strategy. For example, many managers run an offshore and onshore version of the same fund or various funds that follow the same strategy but differ in their fee structure or currency denomination. In essence, when a manager has all funds of the same strategy, it is likely that each fund is a different share class of a single underlying trading strategy.

In my restricted, long/short equity-only sample, I have a total of 16,194 quarterly observations from 602 managers with an average of 2.43 funds per manager.

1.4.2 Definition of Leverage

Hedge funds have several available methods for obtaining effective leverage. The most apparent is borrowing cash, typically through their prime broker, for purchases. Hedge funds can also obtain leverage by using a long-short investment strategy. Every short sale provides cash up front that can be used to purchase additional long positions as long as any risk constraints of the fund are not being violated.

Because of the complicated strategies hedge funds may employ, there are various concepts of leverage that may be used—for a hedge fund there is no unique agreed-upon leverage measure. The primary reason researchers are interested in a fund's leverage is that leverage is a measure of the exposure and risk of the fund on a per-dollar of investor assets under management basis. The use of derivative contracts adds complexities to leverage computation, but even in the case of a fund that has simple long and short positions in equity there are three frequently used leverage measures. First, *gross leverage* is a frequently used measure in the industry and is the total of the absolute position sizes. A fund long \$100M and short \$50m with assets under management of \$50M would have a gross leverage of $150/50 = 3$, often quoted as "three-to-one gross leverage." Gross leverage almost always overestimates the risk of a portfolio because in most cases short positions cancel at least some of the risk of the long positions of the fund. It is therefore the most conservative measure, which is one reason for its popularity. Second, and on the other extreme, is *net leverage*, which is the sum of the signed position values divided by the AUM. In the previous example, the fund has a net leverage of only $(100 - 50)/50 = 1$. Net leverage tends to underestimate the risk of many long-short portfolios, unless the short positions perfectly cancel out the risks of the long positions for a pure arbitrage. Hedge funds with equal long and short assets are not uncommon, especially among those reporting an "equity market neutral" strategy. These funds would have a net leverage of zero, implying—if leverage is used as an exposure measure—that they are as safe as risk-free assets.

A third measure, which is a compromise between the extremes of gross and net leverage, is *long-only* leverage, computed as the value of all long positions of a fund divided by the assets under management. When a fund takes a short position, the cash raised can be used to obtain additional exposure on the long side. Indeed in this case short positions are essentially an alternative method for obtaining leverage. Ang, Gorovyy, and van Inwegen (2011) examine gross, net, and long-only leverage for a sample of hedge funds held by a fund of funds and find that the three tend to move together, with long-only leverage roughly halfway between measured gross and net leverage. While gross leverage treats short positions as only adding risk to a portfolio, and net leverage treats them as if they perfectly reduce it, long-only leverage treats the explicit risk contribution of short positions as zero. They add to the long-only measure of leverage only in that they provide cash that may be used to increase long positions. Since one cannot tell in which cases short positions are used as a hedge and in which they are used to add additional risk/exposure to a portfolio, long-only leverage measurement seems a reasonable compromise. For the duration of this paper I will be primarily interested in the changes in leverage, manager by manager, so the importance of the leverage method is relatively small.

Long-only leverage computation has a significant advantage over other methods: the holdings reported to the SEC through the 13(f) form are taken from a list compiled by the SEC and include only long positions. Thus taking the market value of assets reported to the SEC and dividing by the assets under management as reported to the commercial hedge fund databases, manager by manager, results in a time-series of long-only leverage for each manager. This produces a dataset of hedge fund leverage that covers essentially the universe of reporting hedge funds over the period of high-quality hedge fund data collection (from about January 1994 on).

1.5 Empirical Findings

I construct the estimated leverage ratio and assets for each manager in each quarter. This permits the construction of an index of aggregate hedge fund leverage over time, either by equal-weighting extant managers or weighting them by their beginning-of-quarter assets under management. The resulting indices are plotted Figure 1.1. Several overall trends are apparent in the graph. First, I do not find that there has been a long-term time trend toward higher leverage. Given the dramatic increases in the number and size of hedge funds over time, this calls into question the explanation in Section 1.3.6, that funds lever up as competition drives down profit. As competition in the hedge fund universe has dramatically increased, either per-dollar profitability has not significantly fallen or hedge fund managers have not responded with a long-term trend toward increasing leverage. The graphs illustrate that the time-series variation in aggregate hedge fund leverage is quite large. In both the equal- and asset-weighted indexes, there are fairly dramatic shifts in leverage over time. Comparing the two graphs, the asset-weighted index overall tends to have a lower aggregate leverage level, indicating that managers of large-asset hedge funds may be using less leverage than their smaller competitors. However, I urge caution in interpreting the overall levels. Given the limitations of 13(f) data, I have chosen to focus my analysis on the changes in hedge fund leverage, rather than the overall levels. For example, despite the restrictions I impose on the sample to get equity-only funds, larger hedge funds could have more assets that are not reportable in 13(f) filings and therefore bias the leverage estimate downward relative to their smaller peers.

Vertical lines indicate March 10, 2000 (the day that, numerically, the dot com bubble peaked) and August 2007 (the approximate beginning of the financial crisis). Both of these dates correspond to local maxima in hedge fund leverage. As described by Brunnermeier and Nagel (2004), hedge funds had been aggressively participating in the run up of the dot com bubble and rapidly exited their positions in tech companies during the subsequent crash. At the peak of the housing bubble right before the financial crisis, hedge funds were also investing aggressively, but not to the same degree as they had been during the dot com bubble. In both cases there was a precipitous decline in leverage during the crisis over the course of several years. Again, the decline was more dramatic after the tech bubble burst—perhaps because leverage levels were higher so there was more potential for decline.

My panel data consists of quarterly observations for all matched hedge fund managers passing the filters. This data is summarized by strategy in Panel A of Table 1.1. "Median AUM" refers to the median assets under management reported to the commercial database by the manager's funds, by manager, while "Median Assets" refers to the market value of the assets of those funds as reported in the 13(f) documents. The most common listed strategy is Long/Short Equity with a total of 17,626, though these funds tend to be on the smaller side, with a median reported AUM (per manager) of 225.6 million USD. Event driven funds have a median AUM (by manager) of 536.0. It should be noted that a manager may have funds in more than one strategy. For this table, managers are included in each row in which they have a strategy, so the "All Managers" row contains fewer managers than the sum of the individual strategies.

Because some managers may have funds in more than one strategy, I also create a restricted sample, which consists only of merged manager/quarters for which each manager had funds from a single strategy only. This restricted sample will be used throughout my paper. The restricted sample is summarized in Panel B of Table 1.1. After removing managers with more than one strategy, the restricted sample has 16,194 observations total from 602 distinct managers. The average number of funds per manager in the restricted sample is lower, at 2.43, and the median reported AUM per manager is also somewhat lower, at 200.0 million USD per fund. Removing management companies with funds in multiple strategies removes many of the larger managers.

Table 1.1 also reports the percentage of managers in each strategy that have at least one fund that reports allowing leverage to the commercial database. I consider this dummy unreliable as it is not typically changed over time and may not be interpreted the same way by all managers—for example, selling stocks short and using the cash to purchase long positions would not be considered leverage by many managers. Some 56.3% of all managers run at least one fund that reports being permitted to use leverage.

Table 1.2 gives the total December AUM of the fully merged sample, the subset of the sample that manages equity oriented (equity long/short) funds, and the subset that manages only equity oriented funds. The latter is the most restrictive and cleanest sample.

I utilize a pooled regression methodology to test several of the hypotheses in Section 1.3 relating to the changes in manager leverage. Because the variance of the change in leverage differs from manager to manager, pooling observations across managers and time and using the naive standard error estimate would result in biased statistics. For each pooled regression, I therefore impose clustering of standard errors by manager. This permits the estimated variance of errors to differ by manager and produces consistent t-statistics. See Petersen (2009) for more details.

Table 1.3 reports pooled regression results of hedge fund leverage on some variables that will be used as controls in later parts of the paper. Hedge fund leverage overall is significantly negatively autocorrelated. In specification 3, the AR(1) coefficient is -0.134, with a t-statistic (clustering by manager) of -6.51. I also estimate the AR(1) coefficient for each manager individually, for any manager with more than 8 consecutive quarters of data, and plot the kernel density of the resulting coefficients in Figure 1.3. The mean AR(1) coefficient for these managers is -0.120. The t-statistic of this mean is -6.58, so I reject the hypothesis that leverage is not negatively serially correlated at the quarterly frequency. I therefore include lagged leverage as a control in all leverage regressions.

As a first step toward understanding whether hedge fund managers can use leverage to time their fund returns, I examine the performance of managers who have recently increased or decreased leverage significantly. I form portfolios using an approach similar to that of Jegadeesh and Titman (1993). In each quarter I rank managers into quintiles by percentage change in leverage. I then form subportfolios consisting of managers in the top and bottom quintiles in each quarter. The subportfolios begin the quarter after the ranking and are held for 1, 2, or 3 quarters. Managers within subportfolios are equal-weighted and subportfolio returns are equal-weighted. Note that because of 13(f) and commercial database reporting timings, these portfolios may not be implementable to real investors—they are constructed to illustrate the horizon over which leverage boosts performance.

Figure 1.4 shows the cumulative returns over the sample of the increasing and decreasing portfolios for each holding period. Using a single quarter holding period, the increasing leverage portfolio significantly outperforms the decreasing leverage portfolio. The quarterly performance difference is 1.196% (or 4.869% annualized), which is statistically significant with a t-statistic of 2.25. Increasing the holding period to two quarters, the effect diminishes to 0.600% with a t-statistic of 1.66. Increasing the holding period again to three quarters, the average spread becomes an insignificant 0.53. Using this raw-return approach, managers that increase their leverage outperform those that decrease, but the effect is short lived.

In Table 1.4 I test Hypothesis 1.1: that changes in leverage can be effectively timed to coincide with times of positive concurrent or subsequent risk-adjusted returns. I again run a pooled regression of quarterly asset-weighted returns by manager on contemporaneous and lagged quarterly changes in leverage. I also include standard hedge fund risk factors as proposed in Fung and Hsieh (2004), so the coefficients can be interpreted as effects on risk-adjusted performance. The regression equation is

$$r_{i,t} = \alpha + \beta_1 \Delta\text{Leverage}_{i,t} + \beta_2 \Delta\text{Leverage}_{i,t-1} + r_{i,t-1} + \Phi X_t + \epsilon_{i,t} \tag{3}$$

where $r_{i,t}$ is the asset-weighted return for manager i over quarter t. $\Delta\text{Leverage}_{i,t}$ is the change in computed leverage for manager i between the end of quarter $t-1$ and the end of quarter t and $\Delta\text{Leverage}_{i,t-1}$ is the computed leverage level for manager i at the end of quarter $t-1$. Lagged manager return, $r_{i,t-1}$ is included as a control. X_t contains standard hedge fund risk factors: excess market return (mkt_t), SMB size factor ($size$), change in ten-year treasury yields ($bond$), change in the BAA - ten-year treasury spread ($credit$), and hedge fund mimicking factors for bonds ($PTFSBD$), foreign exchange ($PTFSFX$), and commodities ($PTFSCOM$). After controlling for risk factors, the effect of contemporaneous change in leverage on returns is negative. In specification (3), the coefficient on contemporaneous change in leverage is -0.013 (t-statistic=-4.20) and that of the lagged change in leverage is -0.002 (t-statistic=-0.65). The sign of these coefficients is negative, indicating that, in fact, changes in hedge fund leverage typically harm performance in a risk-adjusted sense.

Hypothesis 1.2 suggests that hedge fund managers, even if they are unable to time profitable returns, may attempt to use leverage to benefit investors by levering up during times of valuable and large investment opportunities. If this is the case (again, even if those investments do not end up producing the expected returns), funds will tend to increase their portfolio concentration at the same time as their leverage increases.

First, we construct a Herfindahl-style measure of concentration for each portfolio, similar to what is done in Goetzmann and Kumar (2008)

$$\text{HHI}_{i,t} = \sum_{j=1}^{N} w_{i,t,j}^2 \tag{4}$$

where $w_{i,t,j}$ is the weight to asset j in manager i's portfolio in quarter t. Higher Herfindahl score computed this way means a higher concentration in a few assets. I also identify the largest weight to a single asset in each manager's portfolio in each quarter as an alternative measure of portfolio concentration. The pooled regression equation is then

$$\Delta\text{Leverage}_{i,t} = \alpha + \beta_1 \Delta\text{HHI}_{i,t} + \beta_2 \Delta\text{MaxWt}_{i,t} + \beta_3 \Delta\text{Leverage}_{i,t-1} + \beta_4 r_{t-1}^m + \epsilon_{i,t} \tag{5}$$

The results are reported in Table 1.6. Changes in both measures of concentration are negatively related to changes in manager leverage. In specification (3) the coefficients of ΔHHI and MaxWt are -0.921 and -0.142, respectively. T-statistics are -1.94 and -3.12. These are contemporaneous changes and indicate that as hedge funds lever up, they increase diversification, rather than loading up on a few assets. Alternatively, as funds lever down, they may liquidate smaller, diversified positions while holding core assets. In unreported tests, I find that the coefficient for the positive leverage change data is statistically indistinguishable from that of the negative leverage change.

In Table 1.7 I test Hypothesis 1.3: that changes in market volatility (forward and backward looking) encourage decreased leverage ratios in hedge funds because of risk controls. I do a pooled regression of the change in leverage over all quarters for all available managers on changes in the volatility environment at the time. The regression equation is

$$\Delta\text{Leverage}_{i,t} = \alpha + \beta_1 \Delta VIX_t + \beta_2 \Delta SPVol_t + \beta_3 \Delta\text{Leverage}_{i,t-1} + \beta_4 r_{t-1}^m + \beta_5 FundSize_i + \epsilon_{i,t} \tag{6}$$

Where $\Delta\text{Leverage}_{i,t}$ is the change in computed leverage for manager i between the end of quarter $t-1$ and the end of quarter t. Similarly ΔVIX_t is the change in the CBOE VIX index (forward looking investor expectations of volatility) and ΔSpxVol_t is the change in realized S&P500 volatility (computed over the previous 3 months using daily returns). I also include manager and market controls: lagged change in leverage and lagged market returns as well as a control for the mean assets under management at the fund. Both the VIX index and 3-month realized S&P volatility are significantly negative in every specification. In specification (3) a change of 1% in the VIX index leads to a -0.442 change in leverage overall (t-statistic=-5.49), while a similar increase in realized S&P volatility leads to a change of -0.273 in leverage (t-statistic=-3.70). In short, I find support for Hypothesis 1.3 using this proxy for market risk/volatility.

I also test Hypothesis 1.4 by regressing changes in hedge fund leverage on changes in interest rates in Table 1.8. I use changes in three rates, in particular: the risk free rate, the TED spread, and the credit spread. The regression equation is

$$\Delta\text{Leverage}_{i,t} = \alpha+\beta_1\Delta\text{Credit}_t+\beta_3\Delta\text{TED}_t+\beta_3\Delta\text{3MoRate}_t+\beta_4\Delta\text{Leverage}_{i,t-1}+\beta_5 r^m_{t-1}+\beta_6 FundSize_i+\epsilon_{i,t}, \tag{7}$$

where $\Delta\text{Leverage}_{i,t}$ is the change in computed leverage for manager i between the end of quarter $t-1$ and the end of quarter t. Similarly ΔCredit_t is the change in the Moodys BAA-AAA credit spread indices published by the Fed, ΔTED_t is the change in the TED (3-month libor minus 3-month treasury) spread, and $\Delta\text{3MoRate}_t$ is the quarterly change in the 3 month treasury rate. I also include lagged changes in leverage and lagged market returns as controls. I use changes in the three month treasury bill rate as a stand-in for the risk free rate, changes in the TED spread for changes in the part of rates born by relatively safe private institutions, and Moody's Seasoned BAA minus AAA Corporate bond yield series as a proxy for the price of default risk for somewhat higher risk institutions. When the rates on BAA bonds is relatively high compared to AAA, investors require a high premium from higher risk borrowers. This will increase the cost of borrowing for hedge funds and coincide with times of relative difficulty in obtaining financing. Changes in both risk measures (the

TED and credit spreads) are negatively related to changes in overall hedge fund leverage (in specification 4, estimates are -4.462 and -8.198, t-statistics are -2.80 and -4.75, respectively). This is consistent with the idea that difficulty and costs of borrowing are binding constraints for hedge funds and cause them to react by decreasing leverage. Interestingly, the risk-free rate actually enters positively here with a coefficient of 2.159 (t-statistic=2.14). The last result is also consistent with the findings in Ang, Gorovyy, and van Inwegen (2011), that hedge fund leverage is negatively related to default risk premium proxies but positively related in some specifications to the libor rate. Changes in the risk-free rate tend to be closely related to monetary policy, and the relation between this policy and the economic conditions governing hedge fund borrowing is uncertain. That is, the risk-free rate is a measure of a certain part of the cost of borrowing, but it may not be a good measure of the difficulty of borrowing (or risk constraints) of a risky investor like a hedge fund.

I also study Hypothesis 1.5, that flows mechanically cause changes in leverage and the managers may quickly trade to return the fund's leverage to its previous level. First I look at the effect of relatively large inflows or outflows (exceeding 10% of the fund's AUM in absolute value) in event space. The first graph in Figure 1.5 illustrates the cumulative changes in leverage for funds having a large inflow during the quarter from 0 to 1. Specifically, I include funds with inflows in the quarter that exceed 10% of the quarter begin AUM of the fund. There is a strong within-quarter decrease in leverage, indicating that as inflows increase the investor capital, the manager does not purchase enough assets to return the hedge funds to the previous leverage level. Moreover, during subsequent quarters the manager does not trade to undo the induced leverage perturbation, but actually continues to decrease leverage. Among explanations for this behavior might be the effect of capacity constraints on the desired positions of the funds or a reduction in desired risk in the presence of new investors, who might be likely to pull out during a sharp downturn. Alternatively, if hedge fund managers are not actively pursuing leverage targets, then changes to fund leverage induced by flows may not induce corrective action. Similarly, large outflows (more than 10%

of the fund's assets lost) mechanically increase leverage and leverage continues to increase more gradually for several quarters after the shock. In both cases an alternative explanation for the continued trend away from the original leverage level would be that the flows continue after the initial shock (positive or negative) for several quarters. This would make sense if share restrictions limited some investor flows. Overall, the event study results suggest that managers do not rapidly trade to rebalance the portfolio to the original leverage level after a large inflow or outflow.

Table 1.9 examines Hypothesis 1.5 in a regression context. Quarterly changes in leverage are regressed on contemporaneous and lagged flows to the fund. The regression equation is

$$\Delta\text{Leverage}_{i,t} = \alpha + \beta_1 \text{Flow}_{i,t} + \beta_2 \text{Flow}_{i,t-1} + \beta_3 \Delta\text{Leverage}_{i,t-1} + \beta_4 r^m_{t-1} + \beta_5 FundSize_i + \epsilon_{i,t} \quad (8)$$

Flows are computed fund-by-fund at monthly frequency in dollar terms

$$\text{DollarFlow}_s = AUM_s - AUM_{s-1} \cdot (1 + r_{j,s-1,s}) \quad (9)$$

where s is a month and AUM_s is measured in dollars. This dollar flow is then summed for all funds within a manager's control and over months within a quarter to obtain the quarterly dollar flow to the manager. The result is then divided by the manager's total AUM at the end of the previous quarter to get a percentage flow. I find that, individually, both contemporaneous and lagged flows are negatively related to leverage decisions. A unit change in contemporaneous (percentage) flows yields a -0.654 change in average leverage, which is significant with a t-statistic of -9.01 in specification (3). After accounting for contemporaneous flows, the negative coefficient of lagged flows is no longer statistically significant. These results are consistent with those of the event study using only large flows: flows do mechanically change leverage and managers do not trade back to the original leverage level, either within the quarter or in subsequent quarters. If managers traded back to the original target, one would expect the coefficient on lagged flows to be positive and significant.

I also test Hypothesis 1.6, that entry of competing hedge funds drives an increase in leverage due to managers attempting to maintain numerically high returns. This would be a indication that hedge fund managers are using leverage in a way that harms investors. Table 1.10 reports a regression of changes in leverage on the lagged number of inceptions in the long/short equity strategy (normalized by the number of funds in the strategy) for various lags. The regression equation is

$$\Delta \text{Leverage}_{i,t} = \alpha + \sum_k \beta_k \text{Inceptions}_{t-k} + \phi_1 \Delta \text{Leverage}_{i,t-1} + \phi_2 r^m_{t-1} + \phi_3 FundSize_i + \epsilon_{i,t} \quad (10)$$

where $r_{t-1,t-k}$ is the return of the funds managed by manager i over the period from $t-k$ to $t-1$, inclusive. Inceptions_{t-k} is the percentage of funds at the end of quarter $t-k$ that are new as of that quarter. Past increases in Inceptions represents increasing entry and competition in a hedge fund strategy. I also include lagged leverage changes and market returns as controls. I find that when changes in leverage are regressed on entry of new hedge funds within the strategy of the funds managed by each manager, the most significant effect is the longest lag of inceptions (9 months), which has a negative coefficient of -1.379 and a t-statistic of -4.63. Increased entry into the strategy has the effect of depressing leverage. This is also consistent with the observation that the overall leverage across all the funds in the restricted sample (Figure 1.1) has trended downward as the number of funds in the universe has increased dramatically.

Looking at the relation between past manager returns and changes in leverage, I find some support for Hypothesis 1.7. Table 1.11 summarizes pooled regressions of changes in leverage on manager lagged returns. For this exercise, I use 3-month, 6-month, and 9-month returns, lagged by three months. The regression equation is

$$\Delta \text{Leverage}_{i,t} = \alpha + \beta_1 r_{t-1,t-k} + \beta_2 \Delta \text{Leverage}_{i,t-1} + \beta_3 r^m_{t-1} + \beta_4 FundSize_i + \epsilon_{i,t} \quad (11)$$

where $r_{t-1,t-k}$ is the return of the funds managed by manager i over the period from $t-k$ to $t-1$, inclusive. I include controls for lagged change in leverage and lagged market returns as in other tests. Returns are computed as the asset-weighted returns of the funds managed by each manager. There is a negative relation between aggregated returns over previous quarters and today's change in leverage. In specification (3) a percent change in nine-month returns for a fund (lagged by 3 months) affects the leverage by -0.087, which is statistically meaningful with a t-statistic of -4.95. However, when other explanatory variables are added to the regression reported in Table 1.13, the sign of the coefficient on 9-month returns flips (estimate=0.047,t-statistic=3.50). I therefore cannot conclude that negative returns drive hedge funds to lever up based on this analysis.

Taking the results of these two tables together, along with the fact that overall hedge fund leverage does not seem to be increasing as competition in the hedge fund space intensifies, it appears that hedge funds do not lever up to boost numerical returns as competition drives down per-trade profitability.

1.5.1 Other Findings

My approach to constructing hedge fund leverage provides the additional benefit of allowing a study of the effect of hedge fund holdings on changes in leverage. In particular I investigate whether certain assets are associated with increasing leverage, and whether those assets are persistent (that is, whether funds lever up by purchasing assets that are hot at the moment or whether they always use the same set of assets). To do this, I first determine which assets are associated with increases and decreases in leverage. I rank all hedge fund managers into leverage change deciles in each quarter. For each asset, I then compute the average rank of all funds that hold it in each quarter. This provides a measure of whether a particular asset is associated with high or low leverage or with changes in leverage. Finally, I compute the average leverage properties of the assets in each manager's portfolio by taking the weighted average of the assets' (averaged) ranks, where the weights are the weights of the assets in each manager's portfolio. Managers with a high asset leverage change rank, for example, hold assets that are associated with funds that are increasing leverage in that quarter.

Table 1.12 reports the empirical relation between the average leverage-change rank of the assets in a manager's portfolio and its leverage changes. It should be remembered that there is a mechanical relation between the contemporaneous changes in leverage-change rank and leverage decision. The only case in which there will be no contemporaneous relation is when leverage changes are completely unrelated to assets (funds levering up are no more similar to each other in asset composition than they are to funds that are levering down). In specification (3) indeed the coefficient of the contemporaneous leverage change ranking of assets is positive and significant (0.196, t-statistic=30.42). The coefficient on the lagged leverage change ranking is statistically insignificant. This suggests that assets associated with increases in leverage do not do so persistently.

1.5.2 Robustness: Joint Tests

Tests of several hypotheses in this paper have regressed changes in leverage on either macro or manager-level explanatory variables. One naturally questions whether some some of these variables may proxy for the same underlying factor. To address this concern, in Table 1.13 I include regressions from each test as explanatory variables, to compare with the individual tests performed earlier. Looking at specification (7), the coefficients of most variables of interest are little changed and retain their statistical significance.

The most significant change in coefficient is that of 9-month lagged returns. Individually, in specification (3), this coefficient is negative and significant (-0.087, t-statistic=-4.95). However, with the addition of the other regressors, the sign flips and is significant (0.047, t-statistic=3.50). Referring to the correlation matrix in Table 1.14, this variable is most highly correlated with changes in the TED spread, which is the other variable with inference that is not robust to inclusion of all the other regressors. In the case of 9-month returns, we can conclude little about the relation between returns and subsequent leverage choices; we rejected the associated hypothesis. In light of these results, I also base my inference about the relation between interest rate changes and leverage primarily on the BAA minus AAA credit spread results.

1.6 Conclusions

Hedge funds are all but unique among financial institutions in their flexibility to employ sophisticated, time-varying leverage as an integral part of their trading strategy. Yet they provide little or no transparency about leverage choices, so researchers (to say nothing of investors and regulators) cannot observe their leverage decisions and the motivations for those decisions. Although they may escape most of the scrutiny paid to more conventional financial institutions, they are important players in financial markets. Increasingly, they have the resources to have a nontrivial impact on prices and often act as a counterparty to important players like investment and commercial banks. As the size of the hedge fund universe increases, the size and importance of hedge fund allocation in investor portfolios necessarily rises. The size of hedge fund positions and risk exposures can vary greatly depending on their use of leverage. For this reason, it is important to understand the dynamics and causes of changes in hedge fund leverage as well as the implications of those changes. Strategic changes in hedge fund leverage over time have the potential to augment the fund's risk-adjusted performance, but there are many other factors that could drive leverage changes in potentially harmful ways.

In this paper, I develop and use a new approach to constructing hedge fund leverage by aggregating information from required 13(f) holdings filings and reported assets under management as reported by commercial hedge fund databases. This method can be replicated by other researchers and provides a view of the changes in hedge fund leverage representing a large portion of funds in commercial databases throughout the time period spanned by available hedge fund data. I use this rich data to test the hypothesis that changes in hedge fund leverage improve risk-adjusted returns and to examine other hypotheses related to strategic changes in dynamic leverage.

Overall I find that hedge funds do not successfully modify leverage to strategically enhance returns: increases in leverage are not related to enhanced contemporaneous or subsequent risk-adjusted returns. Nor do I find that hedge funds increase concentration when levering up, suggesting that they do not use leverage to take full advantage of apparently profitable trading opportunities. I do find that hedge funds modify leverage consistent with risk targeting behavior: they lever down when market volatility rises and when credit spreads increase. This could be due to internal risk-controls and volatility targeting or it could be the result of pressure applied by the prime brokers' risk monitoring mechanisms.

I also find the surprising result that when hedge fund leverage is mechanically perturbed by fund flows, the resulting change in hedge fund leverage appears to be permanent. This suggests that many hedge funds may not maintain and continuously adjust fund holdings to maintain leverage targets. Instead, the path-dependent history of leverage changes due to flows, performance, market rates, and other factors may play a role in deciding a hedge fund's leverage. The implication for investors is that understanding the history of leverage changes is a critical part of inferring a fund's exposure, which is needed to correctly allocate the rest of a diversified portfolio that includes hedge fund exposure.

Regulators and others concerned with market stability need not blame leverage for increased systemic risk posed by hedge funds, as I find a general downward trend in hedge fund leverage. Additionally, hedge funds tend to reduce leverage after periods of poor performance and after increased competition within their strategy. These facts suggest that hedge fund managers are not, in general, using leverage to maintain a false impression of per-dollar profitability nor taking ever greater risks at their investors' expense. However, in light of the rapidly increasing size of the hedge fund industry in general and the largest individual hedge funds in particular, we cannot exclude the possibility that hedge funds, individually or as a group, pose an increasing threat to market stability.

This is the first paper of which I am aware that uses the methodology of combining SEC 13(f) data and the assets under management reported to commercial databases in order to examine these questions. This methodology provides a richer and broader leverage dataset than has been previously available to researchers.

As a byproduct of this leverage construction method, I also obtain the set of assets invested in by each manager over time. Thus this paper's approach provides fertile ground for future research into the relation between actual hedge fund portfolio holdings and borrowing decisions. The apparent opacity of the hedge fund universe need not be an impediment to understanding the fundamentals of hedge fund borrowing and investing behavior.

References

Acharya, Viral, and S. Viswanathan, 2011, Leverage, moral hazard, and liquidity, *Journal of Finance* 66, 99–138.

Adrian, Tobias, and Hyun Song Shin, 2010, Liquidity and leverage, *Journal of Financial Intermediation* 19, 418 – 437.

Agarwal, Vikas, Naveen Daniel, and Narayan Naik, 2004, Flows, performance, and managerial incentives in hedge funds, in *EFA 2003 Annual Conference*.

Ang, Andrew, Sergiy Gorovyy, and Gregory van Inwegen, 2011, Hedge fund leverage, *Journal of Financial Economics* 102, 102–126.

Brunnermeier, Markus, and Stefan Nagel, 2004, Hedge funds and the technology bubble, *The Journal of Finance* 59, 2013–2040.

Chevalier, Judith, and Glenn Ellison, 1997, Risk taking by mutual funds as a response to incentives, *Journal of Political Economy* 105, 1167–1200.

Dai, John, and Suresh Sundaresan, 2010, Risk management framework for hedge funds: Role of funding and redemption options on leverage, Working Paper, Columbia Business School.

Duffie, Darrell, Chenyang Wang, and Hefei Wang, 2008, Leverage management, Working Paper, Stanford University, University of Chicago.

Eichengreen, Barry, and Bokyeong Park, 2002, Hedge fund leverage before and after the crisis, *Journal of Economic Integration* 17, pp. 1–20.

Fostel, Ana, and John Geanokoplos, 2008, Leverage cycles and the anxious economy, *American Economic Review* 98, 1211–1244.

Fung, William, and David Hsieh, 2004, Hedge fund benchmarks: A risk-based approach, *Financial Analysts Journal* 60, 65–80.

Goetzmann, William, and Alok Kumar, 2008, Equity portfolio diversification, *Review of Finance* 12, 433–463.

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *The Journal of Finance* 48, 65–91.

Liang, Bing, and Liping Qiu, 2013, Hedge fund leverage: 2002–2011, Working paper Presented at 2013 Annual FMA Meetings.

Lo, Andrew, 2008, Hedge funds, systemic risk, and the financial crisis of 2007-2008, Statement to the committee on oversight and government reform U.S. House of Representatives.

McGuire, Patrick, and Kostas Tsatsaronis, 2008, Estimating hedge fund leverage, Working Paper 260 Bank for International Settlements.

Petersen, Michell, 2009, Estimating standard errors in finance panel data sets: Comparing approaches, *Review Of Financial Studies* 22, 435–480.

Schneeweis, Thomas, George Martin, Hossein Kazemi, and Vissilis Kawavas, 2005, The impact of leverage on hedge fund risk and return, *Journal of Altenative Investments* 7, 10–21.

Sirri, Erik, and Peter Tufano, 1998, Costly search and mutual fund flows, *Journal of Finance* 53, 1589–1622.

Stein, Jeremy, 2009, Sophisticated investors and market efficiency, *Journal of Finance* 64, 1517–1548.

Figure 1.1: Leverage over time (equity long/short managers)

A plot of estimated aggregate hedge fund leverage over time. I construct estimated long-only leverage for each hedge fund manager in each quarter and then aggregate these values either by equal-weighting managers or weighting them based on their previous end-of-quarter assets under management. The vertical lines are placed at March 10, 2000 and in August, 2007—important dates that can be considered the beginning of the tech bubble crash and financial crisis.

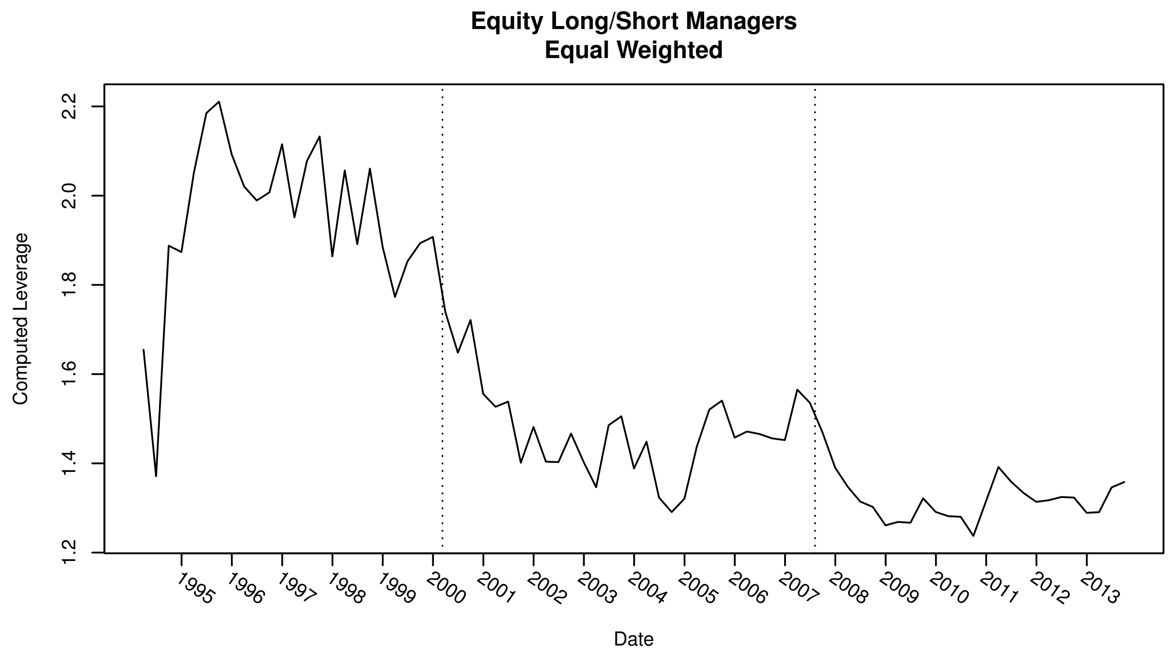

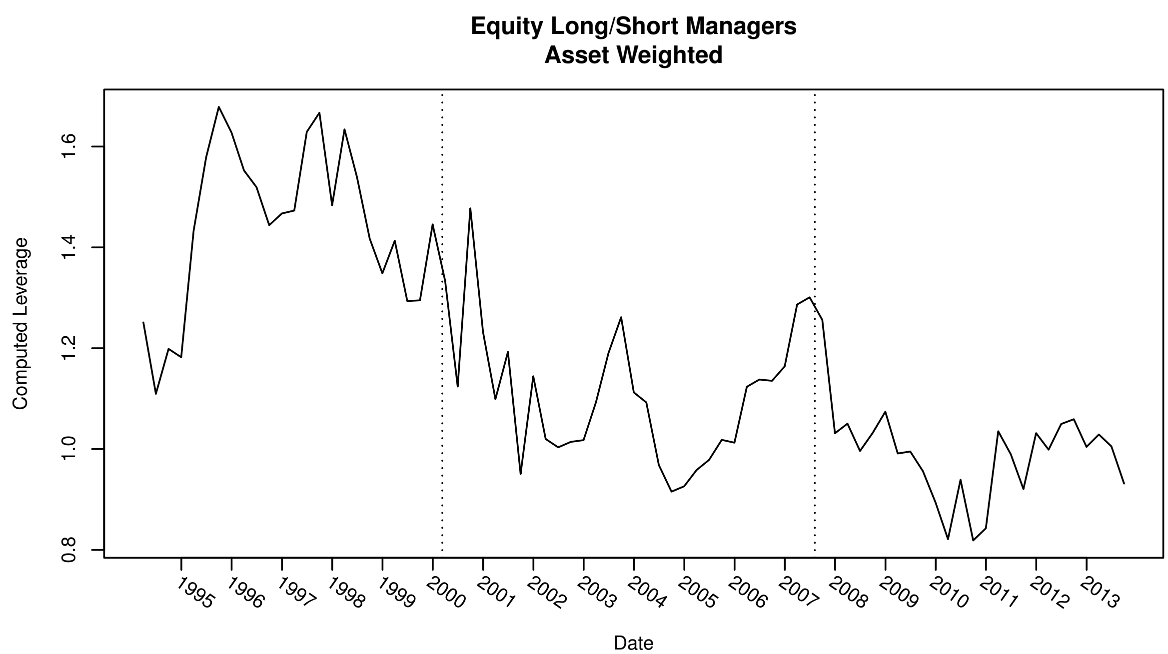

Figure 1.2: Venn diagram of merged manager counts

A Venn diagram illustration of the managers from each commercial database that are found in the 13(f) data. Each illustrated number gives the number of 13(f) managers found in the corresponding database or databases.

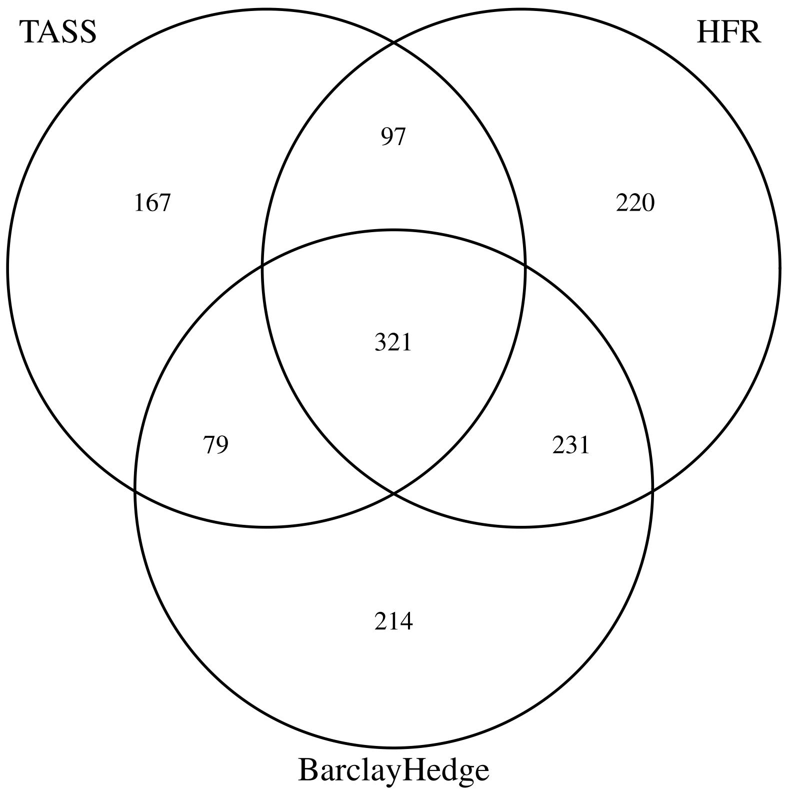

Figure 1.3: Distribution of leverage AR(1) coefficient manager by manager

A kernel density distribution of the AR(1) regression coefficient estimated manager by manager. The regression equation for each manager i is

$$\Delta \text{Leverage}_{i,t} = \alpha_i + \beta_{1,i} \Delta \text{Leverage}_{i,t-1} + \epsilon_{i,t}$$

and the density of the resulting set of $\beta_{1,i}$ is plotted. Managers with more than 8 consecutive leverage values are included. The sample period is from January, 1994 to December, 2013.

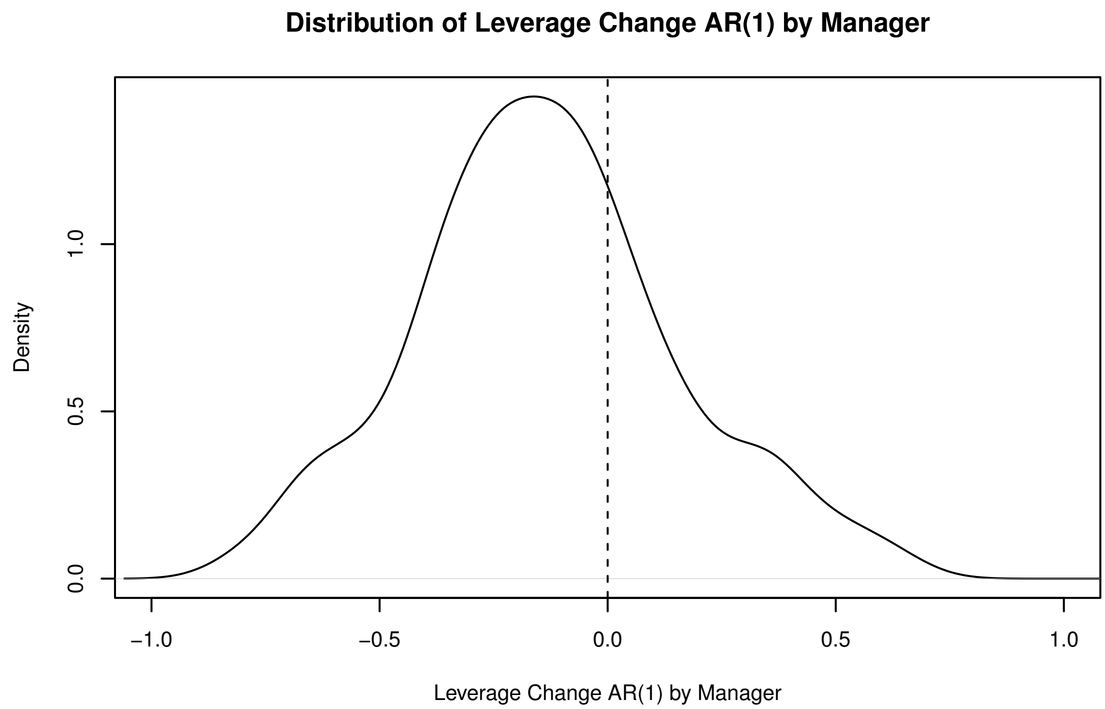

Figure 1.4: Portfolios of managers who increased and decreased leverage

Increasing and decreasing leverage portfolios are formed in a manner similar to Jegadeesh and Titman (1993). Managers are ranked in each quarter by their change in leverage. Portfolios are then formed based on an equal weighting of the returns from managers in quintiles 1 and 5 for each quarter, beginning the quarter after ranking and holding for one or more quarters. The resulting portfolios are then equal weighted for each quarter to give the final portfolio return. The sample period is from January, 1994 to December, 2013.

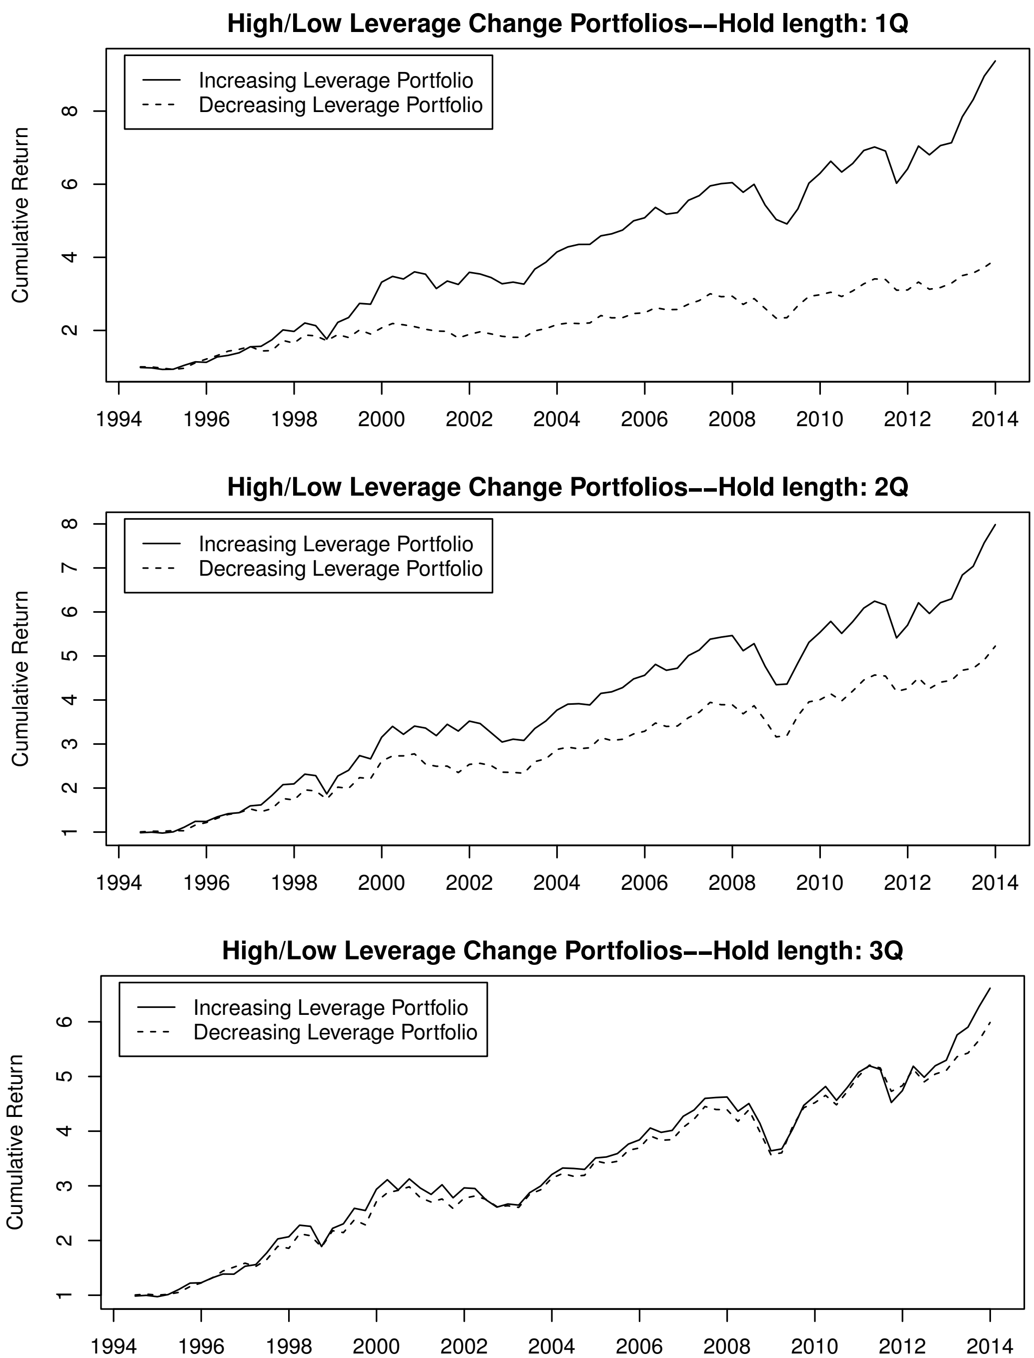

Figure 1.5: Effect of Large Flows on Leverage (LS equity managers only)

Event-time graphs of cumulative changes in percentage leverage after a large inflow or outflow (larger than 10% of the existing AUM of the fund in absolute value).

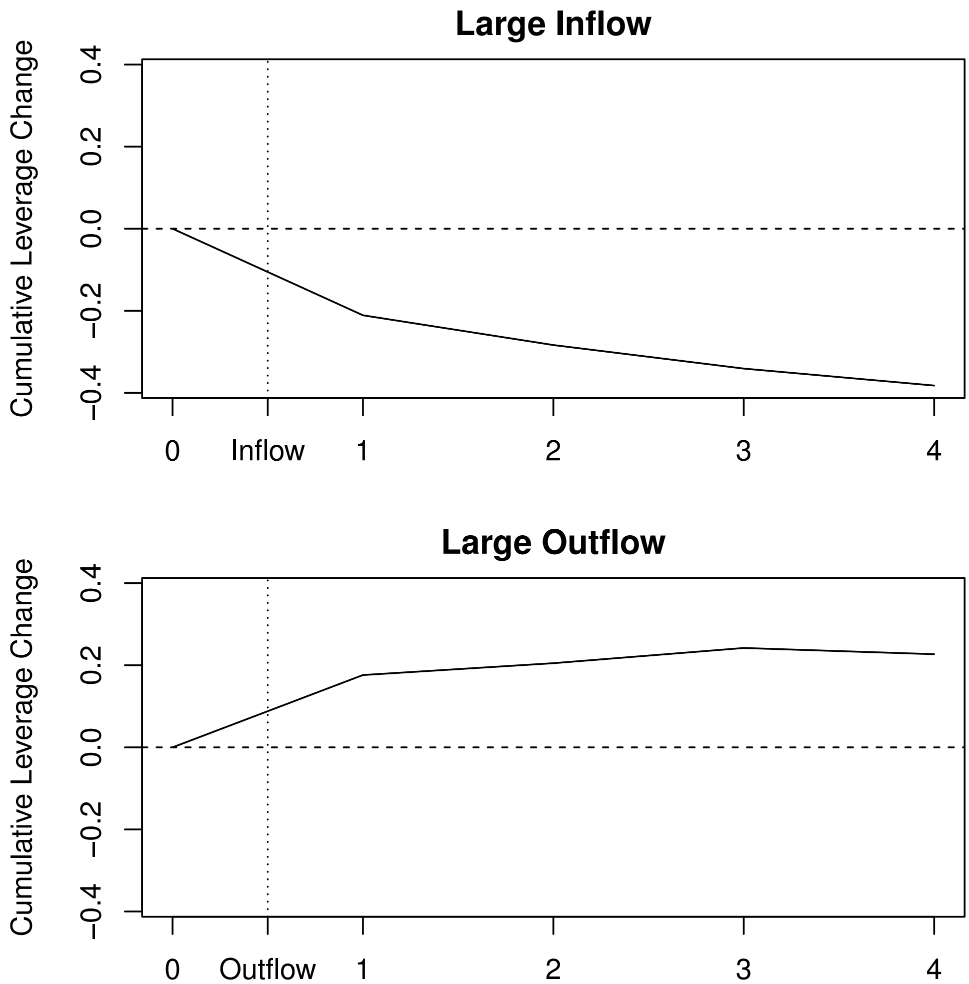

Table 1.1: Summary of managers in the merged sample by strategy.

I construct estimated long-only leverage for each hedge fund manager in each quarter. Note that the observations and number of managers in the "All Managers" row is less than the sum of the individual strategies because some managers manage funds that report two different strategies. In Panel A manager is included in a strategy classification if it manages any funds in the strategy. In Panel B it is included only if all funds it manages are in the strategy. Observations are quarterly by manager. The percent reporting use of leverage is the percentage of managers with at least one fund reporting to the commercial database that they employ leverage. The sample period is from January, 1994 to December, 2013.

Panel A: Full Merged Sample (Some Managers have Multiple Strategies)

	Obs	Managers	Median AUM (Mil)	Median Assets (Mil)	% Reporting Leverage
Long/Short Equity	17,626	728	225.6	321.0	60.3
Event Driven	4,156	172	536.0	298.2	67.7
Other	10,487	558	233.2	283.6	47.4
All Managers	31,395	1,322	245.4	301.9	56.3

Panel B: Restricted Sample (Single Strategy per Manager)

	Obs	Managers	Median AUM (Mil)	Median Assets (Mil)	% Reporting Leverage
Long/Short Equity	13,807	602	186.7	274.5	54.7
Event Driven	2,387	116	385.6	255.9	54.2
All Managers	16,194	713	200.0	271.9	54.6

Table 1.2: Assets under management in merged sample.

This table summarizes the total reported assets under management (AUM) for managers in the fully merged sample, which includes funds from HFR, BarclayHedge, and Tass for which I found a match in the 13(f) filing data. AUM is taken from managers reporting a valid AUM in December of the given year. Equity oriented funds are those reporting a strategy classification similar to equity long/short. Single strategy managers are the subset of equity oriented managers that manage only equity long/short funds. AUM are reported in billions of US dollars.

	Full Sample	Equity Oriented	Single Strategy
1994	26.39	14.63	11.84
1995	41.60	25.72	17.95
1996	55.45	33.95	22.26
1997	81.38	48.30	30.51
1998	88.20	52.22	33.69
1999	113.86	74.79	47.99
2000	142.47	99.55	63.18
2001	181.10	118.97	76.92
2002	194.75	122.07	72.92
2003	281.38	169.28	101.56
2004	447.05	308.08	164.17
2005	588.19	412.22	199.64
2006	894.82	565.61	310.64
2007	1,189.17	777.22	374.83
2008	856.68	474.82	223.17
2009	810.29	516.46	196.61
2010	945.71	592.68	170.57
2011	960.13	563.41	146.10
2012	1,057.24	596.79	130.60
2013	1,174.04	691.70	196.74

Table 1.3: Regression of change in leverage on control variables.

Results of a pooled regression of quarterly changes in leverage by manager on the control variables used in other tables. The regression equation is

$$\Delta\text{Leverage}_{i,t} = \alpha + \beta_1 \Delta\text{Leverage}_{i,t-1} + \beta_2 r^m_{t-1} + \epsilon_{i,t}$$

The sample period is from January, 1994 to December, 2013. Standard errors are clustered by manager and the resulting T-statistics are reported in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	(1)	(2)	(3)
Intercept	-0.003	-0.010***	-0.003
	(-0.78)	(-2.59)	(-0.62)
Lag Change in Leverage	-0.131***		-0.134***
	(-6.49)		(-6.51)
Lag Quarterly SP500 Return			0.117**
			(2.11)
Fund AUM (B)		-0.001	-0.003
		(-0.35)	(-1.35)
R^2	1.8%	0.0%	1.8%
Obs.	8,443	9,076	8,443

Table 1.4: Regression analysis of risk-adjusted manager performance on changes in leverage (LS equity managers only).

Results of a pooled regression of quarterly manager returns on changes in leverage and hedge fund risk measures. The regression equation is

$$r_{i,t} = \alpha + \beta_1 \Delta\text{Leverage}_{i,t} + \beta_2 \Delta\text{Leverage}_{i,t-1} + r_{i,t-1} + \Phi X_t + \epsilon_{i,t}$$

where $r_{i,t}$ is the asset-weighted return for manager i over quarter t. $\Delta\text{Leverage}_{i,t}$ is the change in computed leverage for manager i between the end of quarter $t-1$ and the end of quarter t and $\text{Leverage}_{i,t-1}$ is the computed leverage level for manager i at the end of quarter $t-1$. Lagged manager return, $\text{r}_{i,t-1}$ is included as a control. X_t contains standard hedge fund risk factors: excess market return (mkt_t), SMB size factor ($size$), change in ten-year treasury yields (*bond*), change in the BAA - ten-year treasury spread (*credit*), and hedge fund mimicking factors for bonds ($PTFSBD$), foreign exchange ($PTFSFX$), and commodities ($PTFSCOM$). The sample period is from January, 1994 to December, 2013. Standard errors are clustered by manager and the resulting T-statistics are reported in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	(1)	(2)	(3)
Intercept	0.010***	0.010***	0.010***
	(11.60)	(10.93)	(11.18)
Change in Leverage	-0.014***		-0.013***
	(-4.83)		(-4.20)
Lag Change in Leverage		-0.001	-0.002
		(-0.42)	(-0.65)
Lag Return (VW)	0.063***	0.059***	0.049***
	(4.04)	(3.50)	(2.83)
mkt	0.484***	0.489***	0.493***
	(17.58)	(17.42)	(17.42)
size	0.198***	0.176***	0.190***
	(6.81)	(5.71)	(6.22)
bond	0.431*	0.707***	0.524*
	(1.67)	(2.64)	(1.94)
credit	-1.155**	-0.884*	-1.148**
	(-2.44)	(-1.85)	(-2.32)
PTFSBD	0.003	0.005	0.003
	(0.65)	(1.08)	(0.61)
PTFSFX	0.008***	0.007**	0.008***
	(2.80)	(2.53)	(2.74)
PTFSCOM	-0.035***	-0.037***	-0.034***
	(-8.78)	(-8.60)	(-7.96)
R^2	29.6%	29.6%	30.3%
Obs.	9,062	8,706	8,443

Table 1.5: Correlation of Changes in Leverage and Risk Factors (LS equity managers only).

The correlation matrix of changes in hedge fund leverage and hedge fund risk factors, computed from the pooled leverage sample. $\Delta \text{Lev}_{i,t}$ is the change in computed leverage for manager i between the end of quarter $t-1$ and the end of quarter t. *mkt* is the market excess return, *size* is the (SMB) size factor, *bond* is the change in 10-year treasury rates, *credit* is spread between Moody's Baa and the ten year treasury constant maturity yield. The other three variables are trend following factors: $PTFSBD$ (bond), $PTFSFX$ (currency), $PTFSCOM$ (commodity). The sample period is from January, 1994 to December, 2013.

	ΔLev	mkt	size	bond	credit	PTFSBD	PTFSFX
mkt	12.66						
size	7.74	30.91					
bond	6.72	48.47	31.36				
credit	-9.94	-66.30	-32.63	-70.50			
PTFSBD	-4.76	-40.37	-20.99	-36.39	49.49		
PTFSFX	-6.45	-41.05	-6.03	-38.12	54.80	35.34	
PTFSCOM	-5.75	-35.42	-22.54	-26.41	44.58	46.21	49.40

Table 1.6: Regression of leverage on portfolio concentration measures (LS equity managers only).

Results of a pooled regression of quarterly changes in leverage by manager on contemporaneous measures of portfolio concentration within each management company. The regression equation is

$$\Delta\text{Leverage}_{i,t} = \alpha + \beta_1 \Delta\text{HHI}_{i,t} + \beta_2 \Delta\text{MaxWt}_{i,t} + \beta_3 \Delta\text{Leverage}_{i,t-1} + \beta_4 r^m_{t-1} + \epsilon_{i,t}$$

where $\Delta\text{HHI}_{i,t}$ is the quarterly change in the Herfindahl index applied to the weights of the assets in manager i's portfolio. That is, it is the sum of the squared asset weights in the portfolio. $\Delta\text{MaxWt}_{i,t}$ is the change in the maximum weight to any single asset within the manager's portfolio. $\Delta\text{Leverage}_{i,t-1}$ and lagged S&P returns are included as controls. The sample period is from January, 1994 to December, 2013. Standard errors are clustered by manager and the resulting T-statistics are reported in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	(1)	(2)	(3)
Intercept	-0.002	0.009	0.014**
	(-0.43)	(1.33)	(2.20)
Change in Herfindahl Concentration	-0.854*		-0.921*
	(-1.79)		(-1.94)
Change in Maximium Asset Weight		-0.104**	-0.142***
		(-2.18)	(-3.12)
Lag Change in Leverage	-0.135***	-0.135***	-0.136***
	(-6.57)	(-6.59)	(-6.69)
Lag Quarterly SP500 Return	0.118**	0.114**	0.113**
	(2.15)	(2.05)	(2.07)
Fund AUM (B)	-0.003	-0.003	-0.004
	(-1.44)	(-1.45)	(-1.54)
R^2	2.3%	1.9%	2.5%
Obs.	8,443	8,443	8,443

Table 1.7: Regression analysis of manager change in leverage on risk characteristics (LS equity managers only).

Results of a pooled regression of quarterly change in manager leverage on risk measures

$$\Delta \text{Leverage}_{i,t} = \alpha + \beta_1 \Delta VIX_t + \beta_2 \Delta SPVol_t + \beta_3 \Delta \text{Leverage}_{i,t-1} + \beta_4 r^m_{t-1} + \epsilon_{i,t}$$

Where $\Delta\text{Leverage}_{i,t}$ is the change in computed leverage for manager i between the end of quarter $t-1$ and the end of quarter t. Similarly ΔVIX_t is the change in the CBOE VIX index (forward looking volatility) and $\Delta SpxVol_t$ is the change in realized S&P500 volatility (computed over the previous 3 months using daily returns). I also include lagged ΔLeverage and lagged S&P500 returns as a control. The sample period is from January, 1994 to December, 2013. Standard errors are clustered by manager and the resulting T-statistics are reported in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	(1)	(2)	(3)
Intercept	-0.004	-0.004	-0.004
	(-1.00)	(-0.90)	(-1.04)
Change in CBOE VIX Index	-0.630***		-0.442***
	(-9.20)		(-5.49)
Change in Realized SP500 Volatility		-0.502***	-0.273***
		(-8.02)	(-3.70)
Lag Change in Leverage	-0.136***	-0.137***	-0.137***
	(-6.63)	(-6.71)	(-6.71)
Lag Quarterly SP500 Return	0.222***	0.195***	0.233***
	(4.04)	(3.49)	(4.22)
Fund AUM (B)	-0.003	-0.003	-0.003
	(-1.42)	(-1.38)	(-1.42)
R^2	3.2%	3.0%	3.4%
Obs.	8,443	8,443	8,443

Table 1.8: Regression analysis of changes in leverage on credit constraint proxies (LS equity managers only).

Results of a pooled regression of quarterly leverage changes by hedge fund managers on measures of credit availability

$$\Delta\text{Leverage}_{i,t} = \alpha + \beta_1\Delta\text{Credit}_t + \beta_3\Delta\text{TED}_t + \beta_3\Delta\text{3MoRate}_t + \beta_4\Delta\text{Leverage}_{i,t-1} + \beta_5 r^m_{t-1} + \epsilon_{i,t}$$

Where $\Delta\text{Leverage}_{i,t}$ is the change in computed leverage for manager i between the end of quarter $t-1$ and the end of quarter t. Similarly ΔCredit_t is the change in the Moodys BAA-AAA credit spread indices published by the Fed, ΔTED_t is the change in the TED (3-month libor minus 3-month treasury) spread, and $\Delta\text{3MoRate}_t$ is the quarterly change in the 3 month treasury rate. I also include lagged ΔLeverage and lagged market return as a control. The sample period is from January, 1994 to December, 2013. Standard errors are clustered by manager and the resulting T-statistics are reported in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	(1)	(2)	(3)	(4)
Intercept	-0.002	0.002	0.001	-0.001
	(-0.48)	(0.39)	(0.14)	(-0.21)
Change in BAA-AAA Credit Spread	-9.850***		-8.675***	-8.198***
	(-5.93)		(-5.07)	(-4.75)
Change in TED Spread				-4.462***
				(-2.80)
Change in Three Month Treasury Rate		4.690***	2.979***	2.159**
		(4.76)	(2.95)	(2.14)
Lag Change in Leverage	-0.137***	-0.136***	-0.138***	-0.138***
	(-6.69)	(-6.61)	(-6.74)	(-6.72)
Lag Quarterly SP500 Return	0.093*	0.077	0.071	0.123**
	(1.70)	(1.37)	(1.26)	(2.18)
Fund AUM (B)	-0.003	-0.004*	-0.004	-0.003
	(-1.37)	(-1.74)	(-1.62)	(-1.47)
R^2	2.5%	2.1%	2.6%	2.7%
Obs.	8,443	8,443	8,443	8,443

Table 1.9: Regression analysis of manager leverage on manager flows (LS equity managers only).

Results of a pooled regression of quarterly changes in leverage by manager on flows to the funds in the manager's family. The regression equation is

$$\Delta\text{Leverage}_{i,t} = \alpha + \beta_1\text{Flow}_{i,t} + \beta_2\text{Flow}_{i,t-1} + \beta_3\Delta\text{Leverage}_{i,t-1} + \beta_4 r^m_{t-1} + \epsilon_{i,t}$$

where $\Delta\text{Leverage}_{i,t}$ is the change in leverage for funds managed by manager i over the period from $t-1$ to t. $\text{Flow}_{i,t}$ is the percentage inflows received by the manager's funds over the quarter. Flows are computed fund-by-fund at monthly frequency in dollar terms

$$\text{DollarFlow}_s = AUM_s - AUM_{s-1} \cdot (1 + r_{j,s-1,s})$$

where s is a month and AUM_s is measured in dollars. This dollar flow is then added over funds within a manager's control and over months within a quarter to the quarterly dollar flow to the manager. This is then divided by the manager's total AUM at the end of the previous quarter to get a percentage flow. $\Delta\text{Leverage}_{i,t-1}$ and lagged S&P returns are included as controls. The sample period is from January, 1994 to December, 2013. Standard errors are clustered by manager and the resulting T-statistics are reported in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	(1)	(2)	(3)
Intercept	-0.004	-0.002	-0.004
	(-0.87)	(-0.50)	(-0.86)
Percentage Flow	-0.660***		-0.654***
	(-9.45)		(-9.01)
Lag Percentage Flow		-0.164***	-0.029
		(-4.78)	(-1.02)
Lag Change in Leverage	-0.146***	-0.164***	-0.151***
	(-8.01)	(-7.47)	(-7.50)
Lag Quarterly SP500 Return	0.185***	0.150***	0.188***
	(3.63)	(2.70)	(3.65)
Fund AUM (B)	0.004*	-0.001	0.005**
	(1.82)	(-0.58)	(1.96)
R^2	11.4%	2.5%	11.4%
Obs.	8,333	8,330	8,330

Table 1.10: Regression analysis of manager leverage and competition (LS equity managers only).

Results of a pooled regression of quarterly changes in leverage by manager on the number of entries into the hedge fund strategies of the funds within each management company. Number of entries in a strategy is computed as the sum of hedge fund inceptions in each strategy during the quarter divided by the total number of funds in the strategy at the end of the quarter (i.e., it represents the proportion of funds that are new). The regression equation is

$$\Delta\text{Leverage}_{i,t} = \alpha + \sum_k \beta_k \text{Inceptions}_{t-k} + \phi_1 \Delta\text{Leverage}_{i,t-1} + \phi_2 r^m_{t-1} + \epsilon_{i,t}$$

where $r_{t-1,t-k}$ is the return of the funds managed by manager i over the period from $t-k$ to $t-1$, inclusive. Inceptions_{t-k} is the percentage of funds at the end of quarter $t-k$ that are new as of that quarter. Past increases in Inceptions represents increasing entry and competition in a hedge fund strategy. $\Delta\text{Leverage}_{i,t-1}$ and lagged S&P returns are included as controls. The sample period is from January, 1994 to December, 2013. Standard errors are clustered by manager and the resulting T-statistics are reported in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	(1)	(2)	(3)	(4)	(5)
Intercept	0.009	0.008	-0.004	0.033***	0.020**
	(1.15)	(1.00)	(-0.46)	(3.83)	(2.03)
Strategy Inceptions	-0.377*				-0.059
	(-1.83)				(-0.17)
Lag Strategy Inceptions		-0.328			0.133
		(-1.55)			(0.46)
Lag(2) Strategy Inceptions			0.048		0.649*
			(0.20)		(1.76)
Lag(3) Strategy Inceptions				-1.041***	-1.379***
				(-4.90)	(-4.63)
Lag Change in Leverage	-0.134***	-0.134***	-0.134***	-0.133***	-0.132***
	(-6.53)	(-6.52)	(-6.52)	(-6.52)	(-6.49)
Lag Quarterly SP500 Return	0.127**	0.112**	0.117**	0.104*	0.103*
	(2.29)	(2.01)	(2.11)	(1.87)	(1.83)
Fund AUM (B)	-0.003	-0.003	-0.003	-0.002	-0.002
	(-1.23)	(-1.25)	(-1.36)	(-1.00)	(-1.12)
R^2	1.9%	1.9%	1.8%	2.1%	2.2%
Obs.	8,443	8,443	8,443	8,443	8,443

Table 1.11: Regression analysis of manager leverage and past returns and size (LS equity managers only).

Results of a pooled regression of quarterly changes in leverage by manager on past returns (value-weighted by fund). The regression equation is

$$\Delta\text{Leverage}_{i,t} = \alpha + \beta_1 r_{t-1,t-k} + \beta_2 \Delta\text{Leverage}_{i,t-1} + \beta_3 r^m_{t-1} + \epsilon_{i,t}$$

where $r_{t-1,t-k}$ is the return of the funds managed by manager i over the period from $t-k$ to $t-1$, inclusive. AUM_{t-1} is the assets under management for manager i in the previous quarter. $\Delta\text{Leverage}_{i,t-1}$ is included as a control. The sample period is from January, 1994 to December, 2013. Standard errors are clustered by manager and the resulting T-statistics are reported in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	(1)	(2)	(3)
Intercept	-0.001	0.003	0.004
	(-0.29)	(0.71)	(0.86)
Lag Return (VW)	-0.101		
	(-1.55)		
6 Month Past Return (VW)		-0.150***	
		(-4.76)	
9 Month Past Return (VW)			-0.087***
			(-4.95)
Lag Change in Leverage	-0.135***	-0.136***	-0.135***
	(-6.59)	(-6.60)	(-6.58)
Lag Quarterly SP500 Return	0.176***	0.215***	0.182***
	(2.71)	(3.61)	(3.14)
Fund AUM (B)	-0.003	-0.003	-0.003
	(-1.30)	(-1.16)	(-1.22)
R^2	1.9%	2.1%	2.1%
Obs.	8,443	8,430	8,430

Table 1.12: Regression analysis of manager leverage changes on lagged average asset leverage changes (LS equity managers only).

Results of a pooled regression of quarterly changes in leverage by manager on the average leverage ranking of the assets held by the manager's funds. The regression equation is

$$\Delta\text{Leverage}_{i,t} = \alpha + \beta_1 \Delta\text{LevChangeRank}_{i,t} \beta_2 \Delta\text{LevChangeRank}_{i,t} + \beta_3 \Delta\text{Leverage}_{i,t-1} + \beta_4 r^m_{t-1} + \epsilon_{i,t}$$

To get $\text{LevChangeRank}_{i,t}$ I compute the average leverage change of managers holding each asset at time t. I then rank that average (by asset, quarter) into deciles and find the weighted average of the average asset leverage change rankings for each manager. These rankings correspond to properties of the manager's assets, not of the leverage of the manager's portfolio. $\Delta\text{Leverage}_{i,t-1}$ and lagged S&P returns are included as controls. The sample period is from January, 1994 to December, 2013. Standard errors are clustered by manager and the resulting T-statistics are reported in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	(1)	(2)	(3)
Intercept	-1.088***	-0.002	-1.119***
	(-30.20)	(-0.07)	(-24.63)
Leverage Change Ranking of Assets	0.196***		0.196***
	(30.22)		(30.42)
Lag Leverage Change Ranking of Assets		0.000	0.006
		(-0.02)	(1.58)
Lag Change in Leverage	-0.074***	-0.131***	-0.089***
	(-5.40)	(-5.39)	(-5.28)
Lag Quarterly SP500 Return	0.058	0.112**	0.066*
	(1.57)	(1.97)	(1.79)
Fund AUM (B)	-0.002	-0.003	-0.002
	(-0.74)	(-1.32)	(-0.69)
R^2	47.8%	1.8%	47.8%

Table 1.13: Joint regression analysis of change in leverage on test variables (LS equity managers only).

Results of a pooled regression of quarterly changes in leverage by manager on sets of variables to the funds in the manager's family. The sample period is from January, 1994 to December, 2013. Standard errors are clustered by manager and the resulting *T*-statistics are reported in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Intercept	-0.004	-0.001	0.004	0.033***	-0.004	-1.119***	-1.031***
	(-1.04)	(-0.21)	(0.86)	(3.83)	(-0.86)	(-24.63)	(-22.68)
Change in CBOE VIX Index	-0.442***						-0.110*
	(-5.49)						(-1.87)
Change in Realized SP500 Volatility	-0.273***						-0.423***
	(-3.70)						(-6.26)
Change in BAA-AAA Credit Spread		-8.198***					-3.570**
		(-4.75)					(-2.19)
Change in TED Spread		-4.462***					1.824*
		(-2.80)					(1.72)
Change in Three Month Treasury Rate		2.159**					2.639***
		(2.14)					(3.70)
9 Month Past Return (VW)			-0.087***				0.047***
			(-4.95)				(3.50)
Lag(3) Strategy Inceptions				-1.041***			-0.992***
				(-4.90)			(-6.30)
Percentage Flow					-0.654***		-0.347***
					(-9.01)		(-6.33)
Lag Percentage Flow					-0.029		-0.007
					(-1.02)		(-0.31)
Leverage Change Ranking of Assets						0.196***	0.186***
						(30.42)	(28.79)
Lag Leverage Change Ranking of Assets						0.006	0.006
						(1.58)	(1.31)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	3.4%	2.7%	2.1%	2.1%	11.4%	47.8%	52.2%
Obs.	8,443	8,443	8,430	8,443	8,330	8,414	8,291

Table 1.14: Correlation of Regressors (LS equity managers only).

The correlation matrix of the various regressors used in leverage change regression, computed from the pooled leverage sample.

	ΔVix	ΔSpxVol	ΔCredit	Δ3mRate	ΔTed	9mRet	Lag(3) % Incept	% Flow	Lag % Flow	AvLevRank	AvDeltaLevRank	LagAvLevRank
ΔSpxVol	59.80											
ΔCredit	37.22	69.91										
Δ3mRate	-19.66	-31.12	-28.88									
ΔTed	38.72	42.24	15.18	-23.53								
9mRet	7.26	11.09	3.11	5.62	13.05							
Lag(3) % Incept	-0.18	9.15	9.89	6.97	3.28	6.13						
% Flow	-0.01	-0.75	-0.87	2.25	-0.28	4.58	1.75					
Lag % Flow	1.67	3.18	1.81	-1.67	5.64	5.97	0.01	5.82				
AvLevRank	0.80	0.14	2.31	-3.24	-1.61	-3.16	-12.54	0.60	1.34			
AvDeltaLevRank	-4.98	2.86	3.95	-3.18	-1.84	-8.92	-6.08	-6.75	-2.04	18.24		
LagAvLevRank	-3.17	-2.39	-2.14	-0.66	-5.81	-2.95	-13.36	1.39	0.86	76.20	-6.21	
LagAvDeltaLevRank	1.82	-4.31	0.85	3.93	-7.10	-6.28	-1.07	-1.13	-8.89	9.72	2.63	17.99

Chapter 2: Dollars vs. Sense: Investor Demand, Managerial Skill, and Hedge Fund Startups

Abstract

New hedge funds can be launched either to cater to investor demand or to offer new managerial skills. We hypothesize that skill-driven inceptions deliver better performance. A new hedge fund is identified as investor demand driven when its inception follows high strategy-level returns and inflows, or when it clones existing funds of a same family that have similar investment objective categories. In contrast, new hedge funds launched in low-demand conditions or are not clones of existing funds are more likely to be skill driven. Empirically, skill-driven inceptions outperform demand-driven inceptions by 4% to 5% per year in terms of risk-adjusted return. Our findings suggest that hedge fund startups initiated by managers add more valuable investment opportunities to investors than startups initiated in response to investor demand and that the two types of inceptions may be distinguished beforehand.

2.1 Introduction

The hedge fund industry has been notable for its dramatic growth over the last few decades. It was worth essentially nothing prior to the 1980s but ballooned rapidly to $2.5 trillion in assets under management by 2008. During the financial crisis the industry saw dramatic outflows but has since recovered to around its pre-crisis levels. During this expansion, investor contributions drove large inflows to existing funds and many new funds were created. Although both inflows to existing funds and the inceptions of new funds are of great importance to the structure and growth of the hedge fund industry, the literature has focused only on the former, with a particular eye toward the relation between fund performance and subsequent inflows (or vice versa). By contrast, the literature on inceptions is relatively sparse. This ignorance is surprising in light of the prevalence of hedge fund inceptions in the industry. Most currently operating hedge funds did not even exist a decade ago.

Are investments in new hedge funds essentially substitutes for flows to existing funds, or do new funds contribute new ideas and valuable innovations in management skill? Among new hedge funds, can we identify which inceptions are created to facilitate investor inflows without real innovation and which represent new and potentially valuable ideas? Given the importance of managerial skill in the hedge fund industry, these are some of the central questions investors that researchers must answer. Investors seeking to expand or make new allocations to the hedge fund industry must decide whether to select from existing hedge funds based on their reputation and past performance or to participate in the initialization of new funds. Researchers studying the performance of the hedge fund industry relative to passive investments also need to account for the new investment opportunities offered by startups. We examine these issues by exploring the conditions under which new hedge funds come into being, and by recognizing that barriers to entry and ease of raising capital vary over time by strategy and by hedge fund structure within families. We first notice that although both inception capital and inflows to existing funds can absorb investor capital, there may be reasons for investors to prefer to invest in new funds. Because of diseconomies of scale (e.g., Goetzmann, Ingersoll, and Ross (2003), Getmansky, Lo, and Makarov (2004), Fung, Hsieh, Naik, and Ramadorai (2008), and Teo (2009), etc.), existing funds that grow too large may experience degraded performance. For this reason, investors may seek new funds instead of expanding investments in existing funds, even if the new funds follow similar strategies to existing ones. Existing hedge fund families observing high demand for new funds may create new funds for the purpose of gathering the new assets. If hedge fund inceptions are motivated by the availability of investor capital, we refer to them as *demand driven* (or *investor driven*) inceptions.

Alternatively, inceptions may be initiated by new or existing managers who have identified new investment opportunities or trading strategies. We can think of these skill-driven trading opportunities as positive shocks to the supply of the services offered by the hedge fund industry, increasing the supply of new investment ideas and skill. We refer to inceptions motivated by new talent or ideas as *supply driven* (or *manager driven*) inceptions.

Our key intuition is that these two different types of inceptions may emerge in the hedge fund industry in different periods of time, depending on investors' demand. Indeed, when investors' demand for a certain type of existing hedge fund is high, there are strong incentives and plentiful opportunities for the creation of funds designed solely to absorb the demand (i.e., demand driven). By contrast, when investors' interest is low in general, inceptions are likely to face additional scrutiny as managers of new funds need to convince investors of their value before these funds can raise sufficient capital to go live. Inceptions in the latter circumstance, therefore, are likely to provide new skills or investment opportunities—i.e., they are supply driven.

The above intuition allows us to not only differentiate the role of new hedge funds from that of flows to existing funds—i.e., new skills can be brought to the hedge fund industry via the emergence of new funds in general and supply driven inceptions in particular—but also to make inferences about the potential distribution of performance among new hedge funds. Since supply-driven inceptions are likely to provide new skills and investment opportunities to the hedge fund industry, we hypothesize that this type of inceptions will outperform existing funds. To the extent that demand-driven inceptions aim to absorb investor demand by replicating existing funds, supply-driven inceptions also outperform demand driven inceptions.

We test these hypotheses using the TASS database over the period from 1994 to 2013. Our main findings can be intuitively illustrated by Figure 2.1, which plots the long-term performance (the cumulative abnormal returns), up to 10 years, of both demand-driven and supply-driven inceptions—we discuss the details of how to identify these two different types of inceptions below. The graph clearly shows that supply-driven inceptions outperform demand-driven inceptions in the long run. The performance gap, which can amount to as high as 50% over a 10 year horizon, is economically important.

We rely on three empirical approaches to identify demand driven and supply driven inceptions. In the first approach, we explore whether an inception follows high or low investor demand for investment in its hedge fund strategy category. We use high past strategy inceptions and flows as a proxy for investor demand. When the creation of a new hedge fund follows high (low) strategy demand, we classify it as a demand driven (supply driven) inception.

Our second approach explores the role that hedge fund families play in launching new hedge funds. Importantly, new funds in an existing hedge fund family can closely mimic (a *clone* inception) or be drastically different from existing funds in the family (a *non-clone* inception). A clone inception can happen when fund families facing high investor demand for their existing funds create new yet similar fund to absorb the excess demand. For this reason, a clone inception is likely to be demand driven. By contrast, a non-clone inception is more likely to be supply driven, as the family does not have reputation in the specific strategy yet and needs to convince investors that the new fund can deliver value.

Our final approach combines the above two identification methods and argues that a clone inception is more clearly demand driven when it follows high strategy category demand and a non-clone inception following a period of low strategy demand is especially likely to be supply driven. This combined approach has the most significant power to differentiate demand- and supply-driven inceptions, which we rely on as our main identification to plot the results in Figure 2.1.

We conduct our empirical investigation in three steps. First, we examine how industrial conditions in the previous year affect the likelihood of hedge fund inceptions. This helps us identify candidate proxies for investor demand. We find that strategy category returns and recent number of inceptions are positively associated with the likelihood of subsequent inceptions, suggesting that these variables can be used to predict the time-series of relative investor demand for new funds in each strategy. At the family level, family flows, family inceptions, and the size of the existing assets in the same strategy category of the new fund enhances the likelihood of inceptions. This latter result is consistent with the idea that clone inceptions are more likely to be demand driven.

Next, we examine the performance of demand-driven and supply-driven inceptions based on lagged strategy demand or on family characteristics. That is, we analyze the performance of inceptions based on our first two identification approaches. We find compelling evidence that, in both cases, inceptions in general and supply driven inceptions in particular outperform existing funds. Further, supply driven inceptions outperform demand driven inceptions. For instance, based on the first identification approach of strategy demand, supply-driven inceptions significantly outperform demand-driven inceptions by about 0.295% per month or 3.6% per year. The second identification approach leads to very similar conclusions: supply driven (non-clone) inceptions in existing families significantly outperform demand driven (clone) inceptions by about 1.5% to 2% per year. These results lend initial support to our working hypotheses.

In our third step, we investigate the performance difference between demand- and supply-driven inceptions based on the combination of strategy-level and family-level demand information. Two interesting observations emerge. First, when we double sort funds into groups based on strategy demand and family characteristics, we find that the return spread between strategy-identified demand- and supply-driven inceptions becomes insignificant for funds that are identified as demand driven based on family characteristics. Likewise, the return spread between family-identified demand- and supply-driven inceptions becomes insignificant for funds that are identified as demand driven based on strategy demand. Jointly taking, these observations suggest that an inception is clearly demand (supply) driven only when both strategy and family-level information identify the fund as demand (supply) driven. These results validate our third identification strategy which combines both strategy-level and family-level information to proxy for demand and supply-driven inceptions.

The second interesting observation arises when we examine the performance implication of the integrated identification approach. We find that, similar to the previous cases, supply driven inceptions (i.e., non-clone or new family inceptions following low strategy demand) outperform demand driven inceptions (i.e., clone inceptions following high strategy demand). The performance difference between supply-driven and non-clone demand-driven inceptions, for instance, amounts to 0.375% per month or about 4.6% per year on a risk-adjusted basis, which is higher in terms of economic magnitude than our previous results.

Finally, in addition to the three steps of main analysis, we have also conduct several robustness checks. We find that the findings remain valid when we use alternative approaches to identify the two types of inceptions, when we use alternative time horizon to define the two types of inceptions, and when we control for known bias that may affect the level of reported return of new hedge funds. In brief, our main conclusions that demand-driven inceptions substitute for flows to existing funds and that supply-driven inceptions outperform both existing funds and demand-driven inceptions are well supported by the data.

Our paper makes several contributions to the existing hedge fund literature. It is the first to explore the primary drivers of hedge fund inceptions and their implications for subsequent fund performance in the economic framework of demand and supply. The question of whether hedge fund managers are informed and can deliver superior performance is also at the core of the hedge fund industry. Our results indicate that the distribution of managerial skills in the hedge fund industry could be affected by fund inceptions. Our results also suggest that different types of inceptions are associated with different managerial incentives. Thus, identifying the drivers of inceptions is crucial to understanding the incentives and efficiency of fund managers. As a practical matter, investors seeking innovative and high-performing funds can use our methodology to identify supply-driven funds.

The rest of the paper proceeds as follows. Section II develops the testable hypotheses.

Section III describes the hedge fund data we use in our analysis. Section IV examines the determinants of inception probability. Section V examines the performance of inceptions in times of strong and weak strategy and family demand. Section VI examines the robustness of our methods. Concluding remarks are provided in Section VII.

2.2 Hypothesis Development

An investor seeking exposure to the hedge fund industry has two choices. The first is to invest in an existing fund with a known track record. However, diseconomies of scale are a key characteristic of the type of "arbitrage in expectations" returns expected in the hedge fund industry. Unlike mutual funds, which can scale up significantly without suffering a performance penalty, hedge funds may be reluctant to accept new capital contributions, as discussed in Goetzmann, Ingersoll, and Ross (2003). Otherwise, hedge funds may accept new capital with the possible consequence of needing to seek out second-best alternative investments when the first-choice investment capacity constraints are reached. Either case poses potential problems for a new investor.

The second choice is to turn to a new fund. Hedge fund entry is a common event and a new fund has the potential either to serve as a substitute for an existing hedge fund or to provide new and superior investment opportunities. A new fund, however, may also expose the investor to the risk of having an untested manager and investment strategy that is inferior to existing funds. It is therefore crucial for an investor to know whether hedge fund startups are in general an investment opportunity superior to, on par with, or inferior to existing funds. The average value of hedge fund startups, in turn, depends on whether new fund managers bring profitable trading strategies to the industry, which can be detected from fund performance. This consideration leads to the following hypothesis:

Hypothesis 1 (value-creating inceptions): If new hedge funds, in aggregate, bring valuable new innovations to the investment opportunity set of hedge fund investors, then flows to new hedge funds should outperform flows to existing hedge funds on a risk-adjusted basis.

By contrast, if new hedge funds only aim to provide substitutes for existing funds or to allow unexperienced managers to enter the industry, the performance of new hedge funds will be similar that of existing hedge funds on a risk-adjusted basis.

Next, we notice that investor demand for hedge funds varies across strategy categories over time. In any given period, some categories of assets could appear more attractive to investors than others. Consequently, flows, both to new funds and existing funds, are greater for funds investing in these assets during times of high investor demand.

The impact of investor flow is nontrivial. On one hand, when competition among investors is relatively high, investors' request on performance becomes relatively low (e.g., Berk and Green 2004). As a result, barriers to entry for new funds in these categories are relatively low. The main consideration of these new funds focuses on how to better serve the existing demand, rather than to offer new profitable strategies. When there is no confusion, we refer to inceptions aiming to serve investor demand as demand-driven inceptions.

On the other hand, when investor flows are difficult to come by, managers who seek to start new hedge funds will have a difficult time getting funded. In particular, fund managers who cannot convince investors that they can add value may not successfully start funds when investor demand in their strategy is low. In such times, managers who successfully launch new funds will be motivated more by skill considerations, such as whether their fund will exploit new investment ideas and trading strategies. We refer to inceptions that require additional demonstration of managerial skills as *manager driven* or *supply driven*.

Demand- and supply-initiated inceptions so far defined play very different economic roles and, consequently, deliver very different performance. Demand-driven inceptions are likely to be similar, if not inferior, to existing funds as they are essentially a way to take advantage of favorable conditions in the market for investor capital within a known strategy. By contrast, supply-driven inceptions potentially 0represent new ideas or managers. Inasmuch as these new ideas improve on the existing hedge fund industry, they are likely to represent valuable opportunities for investors. This line of thought leads to the following hypothesis, which can be compared to the alternative view that all inceptions are equally motivated by managerial skills:

Hypothesis 2 (manager and investor driven inceptions): Inceptions during high investor-demand periods act as substitutes for flows to existing funds. Inceptions during low investor-demand periods are more likely to be manager initiated and represent new trading strategies and opportunities. If these trading strategies and opportunities add value, manager-initiated inceptions will outperform investor-initiated inceptions and flows to existing funds.

It should be noted that demand driven inceptions are likely to arise during conditions that are favorable for an investment strategy, but the supply of new managerial talent is not necessarily closely related to these market conditions. As such, during periods of high investor demand, some inceptions will be demand-driven while others will be supply driven. The above hypothesis, however, points out that investor demand nonetheless affects the incentives of managers, as low investor demand imposes an additional burden for managers to demonstrate their skills.

When a hedge fund management company notices high demand for a fund it manages, it has two options. The first is to allow new capital to flow into its existing funds. The second is to create a new fund with the same investment objectives. While the new fund may differ from existing funds in clientele, fee structure, share restrictions, share class, currency, or other fund characteristics, in terms of investment strategy and ideas, it does not represent innovation. We consider flows to these "clone" funds to be substitutes for flows to existing funds.

When a management company faces a shock to its supply of management skill in the form of a new manager or investment idea, it may allow the manager to create a completely new fund with little in common with existing funds in the family. Although the new fund may use the same legal infrastructure as existing funds, its returns and investments are distinct from existing funds. In terms of investment innovation, the new fund is similar to funds that start a new management company. This observation leads to our third hypothesis:

Hypothesis 3 (clones and non-clones): If clone funds are substitutes for existing flows and non-clone funds—whether in an existing family or not—represent new investment strategies and ideas, then funds in new families and dissimilar funds in existing families will outperform clone funds.

Khorana and Servaes (1999) and Aggarwal and Jorion (2008) argue that fund managers have the strongest incentive to deliver alpha in the early stage of their career. The manager of a clone fund does not have these incentives since the fund's reputation depends on an already existing fund with a known track record. Investors in a clone fund are essentially looking for the returns of the existing fund the clone fund mimics.

In order to increase the power of our identification strategy, we interact hypotheses 2 and 3, leading to the following hypothesis:

Hypothesis 4 (demand/supply and clones): If Hypotheses 2 and 3 hold, then clone funds originating during times of high investor demand in a strategy category are the least likely to provide valuable innovations and outperform existing funds. In contrast, non-clone funds in existing families and new family funds launched during times of low demand are most likely to outperform existing funds.

2.3 Data Description

Monthly hedge fund data come from the TASS Database and the sample period extends from 1994 to the end of 2013. We use this database to identify inceptions and compute monthly inflows. Unless otherwise specified, we always report hedge fund return in excess of the risk free rate. Furthermore, we use the seven-factor model from Fung and Hsieh (2001, 2004) to evaluate risk-adjusted return (hereafter: *alpha*). The seven factors are constructed following the instructions from David Hsieh's Hedge Fund Data Library[1].

Hedge funds reporting returns and the time series of AUM to TASS include information about their management company (family) as well as their inception date. They also provide a broad description of their investment strategy by selecting one of the main strategies used by TASS for classification. We refer to this as the fund's strategy category or simply its strategy.

[1] https://faculty.fuqua.duke.edu/~dah7/HFData.htm

Table 2.1 reports the number of funds reporting a valid assets under management (AUM) at the end of each year in each of the 10 hedge fund strategies we examine. These strategies are convertible arbitrage (CA), dedicated short bias (DS), event driven (ED), emerging markets (EM), equity market neutral (EMN), fixed income arbitrage (FI), global macro (GM), long/short equity (LS), managed futures (MF) and multi-strategy(MS). We exclude funds reporting other strategies and those with missing strategies. In particular, we exclude funds of funds to avoid double counting.

There has been a steady increase in the total number of funds in our sample through most of its history, though the number of funds in individual strategies has grown and shrunk significantly over time as funds have exited and either investor demand for new funds or the supply of new managers has languished. We also report the number of distinct managers (families) reported to TASS by our hedge fund sample. Over time the average number of funds per family has increased from 1.38 at the beginning of our sample to 2.97 by the end. The fraction of management families with multiple funds has also increased gradually over our sample, from 0.24 in 1994 to 0.42 at the end of 2013. In all years, more than half of families manage only one fund.

Table 2.2 reports the number of new funds per year as well as the proportion of funds that are new in that year. As the number of live funds in the hedge fund universe has grown over time, the proportion of the universe represented by new funds has decreased from 0.26 in 1994 down to 0.03 in 2013. We also report the total AUM of new funds raised in each year. Because many funds do not report a valid AUM in the month of their inception, we take as “inception AUM” the first nonmissing reported AUM in the first three months of a fund’s life. Funds not reporting a valid AUM in their first three months are excluded from this investigation (and from all analysis involving inception AUM). Inception AUM grew from 1.14 billion USD in 1994 to 12.46 in 2007. Less was raised by new funds during 2008 and 2009, returning to about 12.29 billion in 2011.

We also report new flows to existing funds in each year, partitioned by the age of the fund in the month when flows were received. Flows are computed based on performance an actual or estimated AUM numbers reported to the TASS database:

$$\mathrm{Flow}_{n,t} = \mathrm{AUM}_{n,t} - \mathrm{AUM}_{n,t-1} \cdot \mathrm{R}_{n,t} \tag{1}$$

Here fund flows and AUM are reported in dollars (we convert any AUM reported in another currency to its value in dollars at that time). Later we will use the normalized flow, which is the flow in month t divided by the lagged AUM. $\mathrm{R}_{n,t}$ represents the gross return to fund n in the month between t-1 and t. The time-series correlation between flows to funds less than a year old and new inception AUM is 0.29, while the correlation between new AUM flows and those to older funds is significantly lower: 0.06 for funds between 1 and 5 years old and 0.05 for funds older than five years. In contrast, flows to older funds are relatively highly correlated (0.75 between funds 1 to 5 years old and funds older than 5 years).

Cumulative flows to new funds, young funds, and older funds are graphically illustrated in Figure 2.2. Beginning at the start of the sample in 1994, inception AUM flows steadily climb (AUM flows, by construction, cannot be negative). Flows to first year funds and funds between 1 year and 5 years old climb much more quickly than AUM flows until the financial crisis of 2007, while net flows to funds that are older than 5 years is negative for much of the sample. During the crisis of 2007, funds older than 1 year experienced dramatic outflows. At the same time, funds younger than 1 year experienced only a mild slowdown in growth. This lack of outflow during the crisis may be the result of the lockup period that many funds impose on investors, preventing withdrawals for a period of time after investment in the fund.

Table 2.3 examines the inceptions we will use in our analysis (only funds with 12 valid observations are included) and divides them into groups defined by whether the inception is the first fund in its reported family or is a new inception in an existing family. Additionally, we report how many of the inceptions in existing families are clones of existing funds in their families. A new inception in an existing family is categorized as a clone if it is in the same strategy category (long/short equity, event-driven, etc.) as an existing fund in the family and if the fund has a return correlation with the previously existing fund of 90% or greater. Overall just more than half of the inceptions in our sample (6,156) begin new management companies and nearly half (5,335) are inceptions in existing families. Of the inceptions in existing family, about half (2583) are classed as non-clone funds. Overall we have 11,491 inceptions in our sample.

We construct raw and risk-adjusted performance measures. We measure raw performance by finding the 60 month excess returns of each fund over the risk-free rate. Risk-adjusted returns are computed as the intercept (alpha) from a 60-month regression of fund returns on the 7 hedge fund risk factors used in Fung and Hsieh (2004). The regression equation is

$$\begin{aligned} r_{p,t} = \alpha_p + \beta_{p,1} MKT_t + \beta_{p,2} SMB_t + \beta_{p,3} YLDCHG_t + \beta_{p,4} BAAMTSY_t \\ + \beta_{p,5} PTFSBD_t + \beta_{p,6} PTFSFX_t + \beta_{p,7} PTFSCOM_t + \varepsilon_{p,t}, \end{aligned} \tag{2}$$

where $r_{i,t}$ is the monthly return to fund i in month t, *MKT* is the excess return to the market, *SMB* is the small-minus-big size factor, *YLDCHG* is the monthly change in the ten-year constant maturity yield, *BAAMTSY* is the monthly change in Moody's Baa yield less ten-year treasury constant maturity yield, and the other three are trend following factors created by Fung and Hsieh and available on Hsieh's website: PTFSBD (bond), PTFSFX (currency), PTFSCOM (commodity). The intercept (alpha) from these regressions gives us a measure of the risk-adjusted performance of the portfolio.

2.4 Inception Likelihood

In order to distinguish times of high investor demand from those of low investor demand, we develop proxies for the ease of raising capital for a new fund. These proxies are characteristics of the hedge fund strategy classification and fund family at a point in time. We therefore search for strategy and family variables that are associated with a relatively large number of new hedge funds using a logistic regression methodology on the incidence of a hedge fund inception by date, strategy category, and family.

Table 2.4 reports the results of a logistic regression of all family/year/categories on strategy category and family variables. The dependent variable is set to 1 when a family has an inception in a given quarter and strategy category and zero otherwise. The regression equation is

$$\mathrm{Inception}_{i,j,t} = \Lambda(\beta \times X_{i,t-1} + \psi \times Y_{k,t-1}) + \varepsilon_{i,j,t} \tag{3}$$

Where $\Lambda(\cdot)$ represents the logistic link function and $X_{i,t}$ is a vector of strategy explanatory variables for strategy category i, year t. $Y_{i,t}$ is a vector of family explanatory variables for family j, year t. Strategy return is the average monthly return to an equal-weighted portfolio of funds in a given strategy category in year t-1. Family return is similarly defined. Strategy and family volatility are computed from equal-weighted portfolios of funds over 24 months previous to year t. Strategy and family AUM are the sum of reported assets under management in December of year t-1. Family assets in same Strategy category is the sum of reported assets in strategy category i and family j at the end of year t-1. Normalized strategy inceptions is the number of inceptions by strategy category i in year t-1, normalized by the number of funds in strategy category i at the end of year t-2. Large family open is set to 1 if one of the largest 8 hedge fund families had an inception in strategy category i in year t-1. Inceptions in family is the count of inceptions in family j in year t-1. In all models, year-fixed effects are included as yearly dummy variables in the regression.

Specifications (1) through (4) include only strategy/family/years for families that existed in the previous two years. For these specifications, the coefficients can be interpreted as influences on the probability of a new fund inception in an existing family, but they exclude hedge fund inceptions that begin new families. For specification (5) we include the effects on the probability of a new fund inception by additionally including the first two years of existence for each family (that is, funds that were excluded from the other specifications because they were too new), but we set all family-level variables to zero for these funds. Additionally we include a dummy variable to distinguish the data from the first four specifications from the additional data. Because of the added new-fund observations, the number of observations in specification (5) is actually higher than that of (2), (3), and (4).

Previous strategy category return and inceptions positively predict inceptions in a family/year/strategy with coefficients of 28.53 and 0.79, respectively (t-statistics=6.44 and 2.97) in specification (4). Flow to the family also predicts inceptions in the family (estimate=0.328, t-statistic=6.16 in specification 4), suggesting that, again, many inceptions are demand-driven. High assets in a given strategy also predict inception supporting the idea that clone funds, in particular, come about because existing funds in the family are large and still drawing more capital from investors. The results from section (4) reveal that there is also an interaction effect: families with high assets in the previous year in categories that had good returns in the previous year are particularly likely to have an inception (estimate=12.18, t-statistic=2.02). In specification (5) we add inceptions that begin new families to the regression and zero out all family variables for these funds. This has the effect of including family/quarter/categories for the first couple of years of the families' existence. Broadly speaking, the results are consistent with the specifications that include only inceptions in existing families. Strategy return, in particular, is highly predictive of inceptions in this specification (estimate=36.19, t-statistic=14.67). In specification (5), the strategy flow coefficient actually becomes negative (estimate=-11.07, t=statistic=-6.3). Taken together, family and strategy indicators of inceptions, flows, and returns positively predict future inceptions, indicating that these factors create a high investor demand environment in which initiating a fund is relatively easy.

2.5 Inception Performance

Before categorizing inceptions according to whether they are demand- or supply-initiated we examine which fund and strategy characteristics are associated with high performance over the first few years of the funds' lives. Factors associated with inceptions that are motivated by managers with new ideas, rather than the need for a fund to absorb investor capital, will be associated with higher fund performance if these new ideas add value. We will examine the relation between these variables and both raw and risk-adjusted performance of the inceptions.

In Tables 2.5 and 2.6 we investigate the strategy category and fund-level variables associated with higher inception performance (over the first 60 month of the fund's life) using a regression methodology. In Table 2.5, we regress 60-month excess returns on strategy inceptions, returns, flows, and volatility as well as a dummy variable for whether the fund is the first in its family and whether it is a non-clone inception in an existing family (clone is the omitted dummy variable). We also include year fixed effects and a control for the size of the inception. Because a number of inceptions do not have an associated AUM, we examine this characteristic by creating a dummy variable that is 1 if the fund is missing the inception AUM and also set Inception AUM to 0 where it is missing. The regression equation is

$$\begin{aligned} \mathrm{Performance}_{i,[1,60]} &= \alpha + \beta_1 \mathrm{NormFlow}_{i,[-36,-1]} \\ &+ \beta_2 \mathrm{Inceptions}_{i,[-12,-1]} + \beta_3 \mathrm{Ret}_{i,[-36,-1]} + \beta_4 \mathrm{Vol}_{i,[-24,-1]} \\ &+ \beta_5 \mathrm{FirstInFamily}_i + \beta_6 \mathrm{NonClone}_i + \beta_7 \mathrm{InceptionAUM}_i \\ &+ \beta_8 \mathrm{MissingInception}_i + \varepsilon_i \end{aligned} \tag{4}$$

Where $\mathrm{Performance}_i$ is the cumulative fund return for fund i over the first 60 months after its inception. $\mathrm{NormFlow}_i$ is the lagged flow to the strategy containing fund i before its inception, normalized by the number of funds in that strategy at the beginning of that period. $\mathrm{Inceptions}_i$ is the lagged number of inceptions in the strategy containing fund i before its inception, normalized by the number of funds in that strategy. Ret_i is the lagged strategy return (equal weighted) for the strategy containing fund i previous to its inception, and Vol_i is the strategy volatility over the two years previous to the inception of fund i. $\mathrm{FirstInFamily}_i$ is a dummy variable set to 1 if fund i is the first fund in its family, $\mathrm{NonClone}_i$ is a dummy set to one if the fund is not the first in its family but it is not a clone fund (that is, it is in a new strategy category than existing funds in the family or has a return correlation below 90% with existing funds in the family). The omitted dummy variable is Clone_i, representing inceptions in existing families that are clones of existing funds in the family (same strategy and correlation higher than 90%). $\mathrm{InceptionAUM}_i$ is the total assets under management in the first reported month of fund i's life (set to 0 if no AUM number is reported within the first three months after inception), and $\mathrm{MissingInception}_i$ is a dummy that is set to 1 when there was no inception AUM reported in the first month (this captures the average effect for funds for which we do not know the inception AUM).

In each specification, the lagged strategy flow (normalized by the strategy AUM) negatively predicts subsequent performance of the inceptions in that strategy (-9.094 in specification 4, t-statistic=-4.55). In addition, in specification (3), past strategy return negatively predicts inception performance with a coefficient of -8.148 and a t-statistic of -1.86, though the coefficient is not significant when other variables are added in specification (4). The negative relation between these variables and subsequent inception performance suggests that inceptions in high demand environments as measured by normalized strategy category inceptions have weaker subsequent returns. In specification (4) we also see that dummies for being a non-clone inception in an existing family and being the first fund in a family positively predict subsequent 60 month returns. In addition, strategy volatility positively predicts future performance of inceptions during that period. This suggests that the period after high volatility in a particular strategy is also a period of low demand for new funds, perhaps because investors view volatile categories as less desirable. As a result, inceptions in categories that recently experienced high volatility may also be more likely to be manager-driven and include valuable new investment skill and ideas.

In Table 2.6 we examine the relation between risk-adjusted performance and the inception characteristics from Table 2.5. Our measure of risk-adjusted performance is the alpha coefficient from a 60-month regression of the returns from hedge fund i on hedge fund risk factors as suggested by Fung and Hsieh (2004). After computing alpha for each fund, we regress the cross-section of alpha coefficients on the variables from Table 2.5 in order to determine the effect of strategy and family characteristic variables as of inception on the risk-adjusted performance of hedge funds over the first 5 years of their life. The regression equation is the same as in Equation 3, where the alpha coefficient is our measure of $Performance_{i.}$

In Specification (4), strategy return significantly and negatively predicts risk-adjusted returns with a coefficient of -6.955 and t-statistic of -2.24. Strategy flow negatively predicts performance in specification (1), but does not achieve statistical significance when other variables are added in specification (4). In addition, dummies for being the first fund in a family and being a non-clone inception in an existing family (clone funds are the omitted dummy) are both positive and significant in each specification. The results of Table 2.6 and Table 2.5, taken together, validate the suggestion that our proxies for strategy demand (returns and flows) as well as family-level demand (new family and non-clone) are both related to subsequent fund performance.

We test Hypothesis 1 directly by constructing portfolios of new and existing hedge funds. Subportfolios are created using a 3-month formation period and 60 month holding period. Any new inception in the 3 month period is included in the inception subportfolio and other funds live during that period are included in the non-inception portfolio. Funds within subportfolios are equal-weighted and monthly returns are computed. The resulting 60 subportfolios are also equal-weighted in each month to create final portfolio returns. These portfolio returns are then regressed on the hedge fund risk factors from Fung and Hsieh (2004).

Table 2.7 shows that both the inceptions and non-inception portfolio have significant alphas (0.362% and 0.240% per month, respectively) and the portfolio formed by going long the inceptions portfolio and short the non-inception portfolio has a significant and positive alpha as well (0.122% per month or 1.46% per year, t-statistic=3.78). This last result supports hypothesis 1, that inceptions, on average, do outperform existing funds for the first 5 years. The implication is that new inceptions bring economically valuable new ideas to the hedge fund universe and investors who take advantage of those inceptions benefit from doing so.

Also in Table 2.7, we test Hypothesis 2 by constructing a measure of investor demand by strategy category. We choose two variables: 36-month flows into a strategy category and 36-month returns in the strategy. In each month we rank the 10 strategies using these two lagged variables. Categories with high rank (7 or greater) are those experiencing high relative investor demand according to that variable and categories with low rank (4 or lower) are experiencing low investor demand according to that variable. To determine which categories have high investor demand overall in month t, we choose those that rank high by both measures in month t. Similarly, categories that rank low by both measures in a particular month are experiencing relatively low investor demand. We then form portfolios of high demand and low demand using the 3 month formation period and 60 month holding period as in Table 2.6 and compute the risk-adjusted return. The results are in the last 3 columns of Table 2.5. The low strategy demand portfolio has a statistically significant alpha of 0.483% monthly, while the high strategy demand portfolio does not have a significant alpha. Further, the monthly spread between the two portfolios is 0.295% per month or 3.54% per year, which is significant (t-statistic=2.03). This result supports Hypothesis 2 in that inceptions in categories experiencing low demand subsequently outperform higher-demand inceptions. If inceptions in low-demand environments are essentially substitutes for flows to existing funds while inceptions in high-demand environments are likely to bring new and valuable managerial skill and investment ideas to the industry, we would indeed expect to see a positive and significant spread here as the new ideas in supply-driven funds cause superior performance in those funds.

In Table 2.8, we examine the effect of family relationships and test whether inceptions that create a new family outperform inceptions in existing families. We decompose inceptions into new family inceptions and existing-family inceptions and forming portfolios similar to the inception and non-inception portfolios. We then estimate Fung-Hsieh seven factor regressions to obtain the risk-adjusted performance measure (alpha) for each portfolio. Grouped this way, both new family inceptions and inceptions in existing families have positive alpha (0.372% and 0.338%, respectively, with t-statistics of 6.06 and 4.55). However, the portfolio formed by the spread between these portfolios does not have a statistically significant alpha. This could be because the existing-family inceptions portfolio contains non-clone inceptions, which we expect to have positive risk-adjusted performance similar to the family inceptions.

Also in Table 2.8, we test Hypothesis 3 by decomposing inceptions within existing families into those that are classified as clones and those that are not. We form clone and non-clone portfolios using the same technique as in Table 2.5 and compute the Fung-Hsieh alphas of each portfolio and the spread between them. The non-clone hedge fund inception portfolio significantly outperforms the clone portfolio. The alpha on the spread is 0.129 per month or 1.55% per year (t-statistic=2.61), providing support for Hypothesis 3. Non-clone funds are likely to provide new ideas that contribute to superior subsequent performance when compared to clone funds. Flows to the latter are essentially equivalent to flows to existing funds. The R^2 in the non-clone and clone specifications is quite high (53.7% and 54.7%) because the hedge fund risk factors included in the regression were designed to explain the variance in the returns of portfolios of hedge funds. On the other hand, they are less effective at explaining the variance of the portfolio that is long non-clone and short clone hedge funds. Most of the factor loadings in the clone and non-clone portfolios are similar, so hedge fund risk factors have little power to explain the differences in the spread.

Figure 2.3 graphically illustrates the performance of new family inceptions, clones, and non-clone inceptions in existing families. Cumulative abnormal returns are plotted from the time of inception up to 120 months after inception. The poorest performing group of inceptions is the clone funds. The other set of inceptions that underperforms the all-inceptions curve is the set of all inceptions in existing funds (this contains both clone and non-clone partitions). Both new family inceptions and non-clone inceptions in new families outperform the all-inceptions line, and their performance is similar. This suggests that new funds in existing families that are not classified as clones are similar to inceptions that create new families in terms of the innovations and profitable opportunities they provide.

We then combine our fund-structure results from studies of Hypothesis 2 and 3 and partition inceptions both by family structure (clone, non-clone in existing family, new family) and inceptions that occurred in high- and low-demand environments. For each of these 6 partitions of inceptions, we form a portfolio using the same 3 month formation period and 60 month holding period as in previous tables. We then regress these portfolio returns (as well as pairwise and corner spreads) on Fung Hsieh factors to get the risk adjusted alpha. We report results from the most interesting of these regressions in Table 2.10 and summarize the matrix of alpha coefficients in Table 2.9. Looking at Table 2.9, we see that the low-high portfolio for new family inceptions has positive and statistically significant alpha (estimate is 0.292 with t-statistic 2.66). For non-clone and clone inceptions the spread is statistically insignificant. Looking the other direction, within high and low groups, the non-clone minus clone and new family minus clone alphas are not significant. One explanation would be that non-clone funds and new families experience only moderate barriers to entry when strategy demand is high, which permits mediocre managers to start new funds in that environment.

The leftmost two entries in Table 2.9 on the bottom row report the alphas of the spread between non-clone funds in low strategy demand periods and clone funds in high demand periods (the corners) as well as new family funds in low strategy demand periods and clone funds in high demand periods. In both cases the alpha are similar and positive (0.375% and 0.380%) and significant (t-stats=2.12 and 2.40). The latter case is illustrated graphically in the cumulative return graph in Figure 2.1 in the introduction. Low-demand new-family funds are the most likely to represent new and profitable ideas and trading strategies, while clone funds in high demand categories are likely to represent a response to investor demand for hedge funds. The sizeable performance gap between manager-initiated and demand-initiated funds in this figure supports Hypothesis 4 in that manager-initiated funds represent innovation and new ideas and these funds outperform demand-initiated funds.

Overall, our results suggest that both our strategy-level measures (past strategy returns and past strategy flows) and family-level measure (whether the inception is a clone or not) of whether an inception is supply-driven are associated with superior subsequent performance. Moreover, inceptions that rank low in both strategy and family measures (supply driven inceptions) outperform inceptions that rank high by these measures (demand driven inceptions) even more dramatically. Our overall interpretation is that supply-driven inceptions contribute significant and valuable new managerial skill to the hedge fund industry when compared to flows to existing funds or demand-driven inceptions.

2.6 Robustness of Our Methodology

Backfill Bias

Because the hedge fund data contained in TASS is voluntarily reported, the potential for biases in the data exist. Of the primary concern to researchers using commercial databases to study hedge funds is the possibility of backfill bias. Many funds decide to submit their performance to TASS after they have been live for some time. At the time that they start reporting, they may optionally add the historical performance for their fund to the database. Because managers want to portray their funds in as positive a light as possible, funds may decide not to include historical performance numbers from their inception to the date at which they started reporting if that performance was no good. Conversely, funds that have reason to be proud of their performance before the date at which they begin listing will include their historical performance. For this reason, available inception performance numbers from the early in funds' lives are likely to have an upward bias as described in Fung and Hsieh (2000). This potential bias is a particular concern for our analysis because since this paper compares the performance of inception portfolios to portfolios of seasoned hedge funds. More generally, much of our analysis is primarily concerned with the performance of funds during the first few years of their lives.

Thankfully TASS also includes a variable designating the date at which the fund started reporting its returns to the database. By dropping out the performance for each fund previous to the date it started reporting, we can obtain a sample that is free of backfill bias; that is, in the backfill-free sample performance numbers were all reported on a monthly basis to TASS as they happened, precluding the reporting flexibility that drives the bias.

In Table 2.11, we construct high and low demand portfolios using the same variables and methods as in Table 2.7 using the bias-free subsample. We perform a regression of the first 60 months of returns from each fund on the Fung Hsieh risk factors to identify the alpha. In the bias-free sample, the high demand portfolio does not have a statistically significant alpha, while the low demand portfolio has an alpha of 0.489% per month or 6.03% per year (t-statistic=5.26). Additionally, the risk-adjusted alpha on the spread between high and low demand portfolios is positive (0.293%) and significant (t=2.00), suggesting that inceptions from low strategy-level demand periods, which are driven by the supply of managerial skill in our model, significantly outperform those from high strategy-level demand periods. This is very close to our result from the full sample and suggests that backfill bias is not a major factor in our portfolio results.

Alternative Proxy Specification

Several of our specifications suggest that past strategy category inceptions (normalized by the number of funds in the strategy) is a good proxy for investor demand. We therefore substitute previous 12-month normalized inceptions at the strategy level for the strategy flows and returns that we used throughout the paper as proxies for investor demand. Again using 3 month formation periods 60 month holding periods we construct returns and compute risk-adjusted alpha for high and low demand portfolios and report the results in Table 2.12. As in Table 2.7, the Fung and Hsieh hedge fund risk factors are very effective at explaining the variance in the returns of both the low and high demand portfolios (R^2 are 52.2% and 54.4%, respectively). Again the risk loadings of the two portfolios are similar, so that the risk factors explain relatively little of the variance in the spread portfolio (R^2 is 13.3% in the spread portfolio. Overall the alpha of the spread portfolio is positive (0.154) and statistically significant (t=2.03), suggesting that inceptions during times of low strategy demand, which are likely to be driven by the supply of manager skill, outperform inceptions from categories experiencing high investor demand. These results are similar to those of Table 2.7. This similarity suggests that our result that fund inceptions during high demand periods underperform those of low demand periods is not sensitive to choice of demand proxy.

Alternative Clone Fund Cutoff

In order to determine whether an inception within an existing family is a clone of a fund within that family, we required that the fund report the same strategy category as the previously existing fund and that the return correlation between the two funds is at least 90% over the period when both funds were live. The choice of 90% is somewhat arbitrary. Indeed, one can easily imagine two funds run by completely different managers in the same strategy category with greater than 90% correlation or perhaps clones in the family that have enough differences in their implementation that their correlation is below 90%. We create an alternative clone/non-clone measure using a lower cutoff threshold of 85% and examine clone and non-clone portfolios using this alternative measure.

Table 2.13 reports results of regressions of clone and non-clone portfolio returns on hedge fund risk factors, including the Fung and Hsieh 7 factors. Subportfolios are created using a 3-month formation period and 60 month holding period. Any inception in an existing family in the 3 month period is included in either the clone or non-clone portfolio, depending on its classification. Funds within subportfolios are equal-weighted and monthly returns are computed. The resulting 60 subportfolios are also equal-weighted in each month to create final portfolio returns. These portfolio returns are then regressed on the hedge fund risk factors from Fung and Hsieh (2004).

In Table 2.13, the risk adjusted alpha of the non-clone portfolio using the alternative classification measure is 0.407% monthly, which is significant with a t-statistic of 5.23. The clone portfolio using this looser classification also is statistically positive (estimate=0.302%, t-statistic=3.72). The portfolio formed by the spread between these two portfolios also has a statistically significant alpha of 0.106% monthly with a t-statistic of 1.97. Loosening the classification criteria to include funds with between 85% and 90% correlation with existing funds in the family reduces the performance gap between clone and non-clone portfolios. This suggests that funds with this relatively lower correlation with existing funds may not be clones. True clones (in the sense that they copy existing funds with little or no difference) are likely among those with the highest correlations with existing funds in the family.

2.7 Conclusions

This paper explores the conditions under which new hedge funds are launched. We propose that proposes that hedge fund inceptions can be motivated either by investor demand for new funds or by managerial, supply side, prerogative. The latter group represents new managers or new ideas from current managers. We classify new hedge funds as manager-motivated or investor demand-motivated using proxies for the economic environment at the time of inception. Specifically we look at whether the strategy category has recently performed well and received large flows relative to other strategy categories. We also look at whether the new fund arises within an existing fund family or starts its own family, and whether funds in existing families are clones of previously existing funds within the family. We find many cases in which new hedge funds are launched in response to investor demand. While these inceptions provide an important channel through which new external capital can be allocated to hedge funds without investing directly in existing funds, they do not represent new and value-enhancing investment opportunities. As judged by performance, our results suggest that demand-initiated inceptions are substitutes for flows into existing funds.

We identify inceptions that are likely to be initiated by managers based on both of our identification criteria. New hedge funds that are the first in their family (or are the first in their strategy category within a family) and new hedge funds that do not follow a periods of high strategy-level demand, as measured by high inflows and returns relative to other categories, can deliver superior performance after the inception dates for 10 years. Manager-initiated inceptions effectively supply new talent to the hedge fund industry and may have a profound impact on the distribution and evolution of skill in that industry.

These findings suggest that there is considerable variety in the degree to which new inceptions bring new and profitable ideas and opportunities to hedge fund investors. Further, they suggest that it may be possible to distinguish new funds that bring genuine innovations to the industry from those that are simply taking advantage of investor's desire to get hedge fund exposure. By choosing the former inceptions, a hedge fund investor can outperform the industry overall.

Hedge funds are unique among investment intermediaries in that their foundation rests on the principle that hedge fund managers possess superior skill and valuable investment ideas. Our approach provides a new way for researchers to understand this central question: how does the distribution of skill in the hedge fund industry evolve and what are its effects on the performance of the hedge fund industry? To do this we must disentangle the effects of investor demand from those of the supply of manager skill to the industry. By isolating inceptions that are likely to be motivated by shocks to the supply of managerial skill, we identify a critical channel through which new ideas enter the industry.

References

Agarwal, R., Jorion, P, 2010. The performance of emerging hedge funds and managers. Journal of Financial Economics 96, 238-256.

Agarwal, V., Daniel, N., Naik, N, 2004. Flows, performance, and managerial incentives in hedge funds. Working paper 04-04, Centre for Financial Research Cologne.

Baquero, G., Verbeek, M, 2007. A portrait of hedge fund investors: Flows, performance, and smart money. Working paper, Erasmus University.

Berk, J., Green, R., 2004. Mutual fund flows and performance in rational markets. Journal of Political Economy 112, 1269-1295.

Chevalier, J., Ellison, G, 1997. Risk taking by mutual funds as a response to incentives. Journal of Political Economy 105, 1167-1200.

Ding, B., Getmansky, M., Liang, B., Wermers, R., 2009. Share restrictions and investor flows in the hedge fund industry. Working paper, SUNY Albany, U Mass at Amherst, University of Maryland.

Fung, W., Hsieh, D., 2000. Performance characteristics of hedge funds and commodity funds: Natural vs. spurious biases. Journal of Financial and Quantitative Analysis 35, 291-307.

Fung, W., Hsieh, D., 2001. The risk of hedge fund strategies: Theory and evidence from trend followers. Review of Financial Studies 14, 313-341.

Fung, W., Hsieh, D., 2004. Hedge fund benchmarks: A risk-based approach. Financial Analysts Journal 60, 65-80.

Fung, W., Hsieh, D., Naik, N., Ramadorai, T., 2008. Hedge funds: Performance, risk and capital formation. Journal of Finance 63, 1777-1803.

Getmansky, M., Lo, A., Makarov, I., 2004. An econometric model of serial correlation and illiquidity in hedge fund returns. Journal of Financial Economics 74, 529-609.

Goetzmann, W., Ingersoll, J., Ross, S., 2003, High-water marks and hedge fund management contracts. The Journal of Finance 58, 1685-1718.

Huang, J., Wei, K., Yan, H., 2007. Participation costs and the sensitivity of fund flows to past performance. Journal of Finance 62, 1273-1311.

Keswani, A., Stolin, D., 2008. Which money is smart? Mutual fund buys and sells of individual and institutional investors. Journal of Finance 63, 85-118.

Jegadeesh, N., Titman, S., 1993. Returns to buying winners and selling losers: Implications for stock market efficiency. Journal of Finance 48, 65-91.

Khorana, A., Servaes, H., 1999. The determinants of mutual fund starts. Review of Financial Studies 12, 1042-1074.

Lynch, A., Musto, D., 2003. How investors interpret past fund returns. Journal of Finance 58, 2033-2058.

Sapp, T., Tiwari, A., 2004. Does stock return momentum explain the smart money effect? Journal of Finance 59, 2605-2622.

Sirri, E., Tufano, P., 1998. Costly search and mutual fund flows. Journal of Finance 53, 1589-1622.

Teo, M., 2009. Does size matter in the hedge fund industry? Working paper, Singapore Management University.

Zheng, L., 1999. Is money smart? A study of mutual fund investors' fund selection ability. Journal of Finance 54, 901-933.

Figure 2.1: Cumulative Abnormal Return after Inception by Investor Demand at the Time of Inception

Cumulative abnormal returns are plotted in event time for all available inceptions, classified by whether investor demand in their strategy was high or low at the time of inception. Strength of investor strategy category demand is measured by two variables: the magnitude of the flows into that strategy during the previous 36 months and the returns in the strategy over the previous 36 months. Inceptions in strategies in the top 4 deciles by both measures are considered high demand while inceptions in the bottom 4 deciles by both measures are considered low demand. Family-level demand is proxied by whether the inception represents a new fund family (high demand) or a clone in an existing family (low demand).

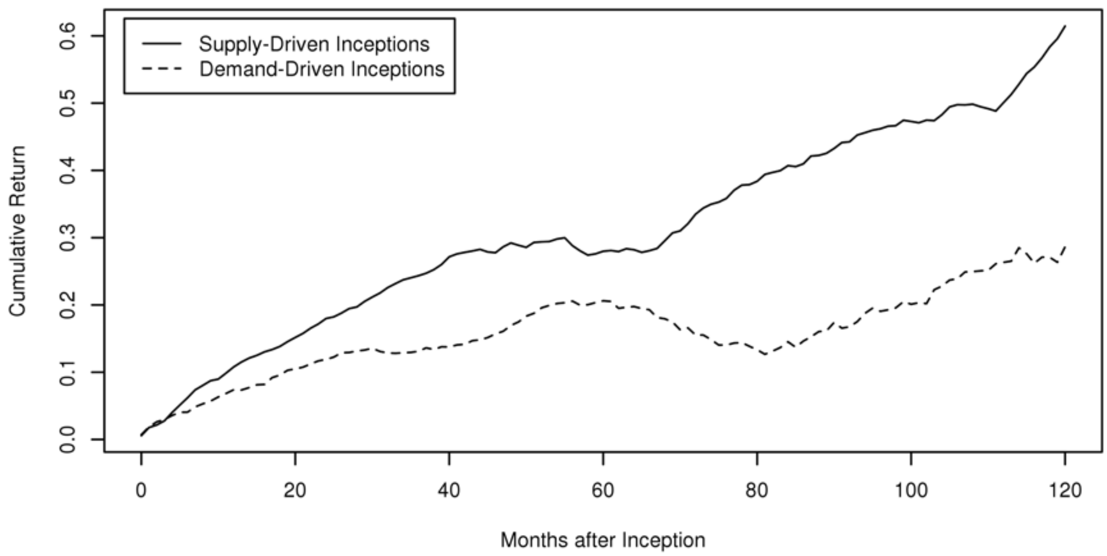

Figure 2.2: Cumulative Flows to Inceptions and Existing Funds

Cumulative flows (since the beginning of the sample) to various types of funds are plotted. Initial flows to new funds are given by AUM of Inceptions, which is the first nonmissing AUM reported to TASS within three months of the fund's inception, or zero if none is available. Flows to first-year funds are any flows to funds that, in the month of the flow, were between the age of 1 and 12 months. Flows to funds 1-5 years are any flows to funds that, in the month of the flow, were between 25 and 60 months old. Flows to funds older than 5 years are any flows to funds that, in the month of the flow, were older than 60 months.

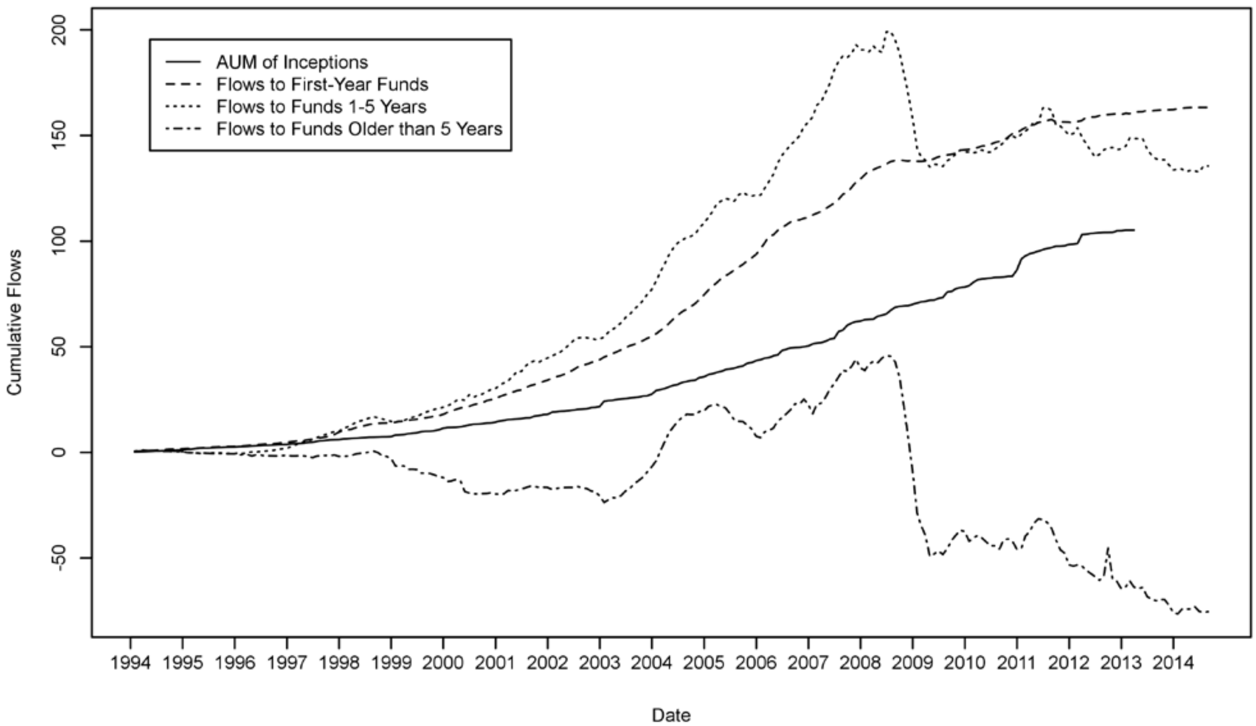

Figure 2.3: Cumulative Abnormal Returns for Inceptions by Type

Cumulative abnormal returns are plotted in event time for inceptions of various types. New family inceptions are fund inceptions that represent the first fund for their reported family. If multiple funds begin on the first date of the family, they are all considered new family inceptions. We classify inceptions in existing families into clone and non-clone funds. Clone funds are those in the same reported strategy as an existing fund in the family and with a return correlation with the previous fund of above 90%. Non-clone funds are inceptions in existing families but in a new strategy or with a return correlation below 90% with existing funds in their family.

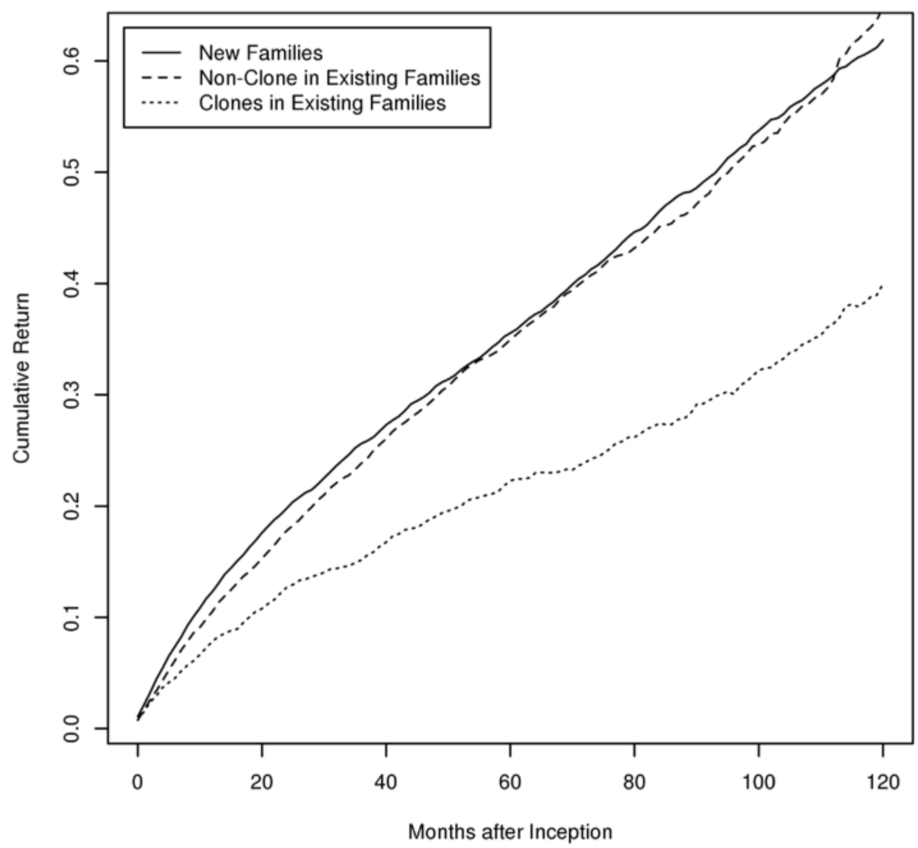

Table 2.1: Fund and Family Counts

This table reports the number of funds in our sample by investment objective strategy category and year. Funds are included if they report a valid AUM in December of the given year. Strategy classifications are convertible arbitrage (CA), dedicated short bias (DS), event driven (ED), emerging markets (EM), equity market neutral (EMN), fixed income arbitrage (FI), global macro (GM), long/short equity (LS), managed futures (mf), and multi-strategy (MS). Excluded are funds without a reported strategy classification in this list (for example, funds of funds). We also report the total number of hedge funds and hedge fund families, as well as the average number of funds per family and fraction of families that have multiple funds in each year. Data are taken from funds reporting a valid AUM to TASS in December of the given year. The sample extends from 1994 through 2013.

	Strategy Category Counts													
Year	CA	DS	ED	EM	EMN	FI	GM	LS	MF	MS	Number of Funds	Number of Families	Funds per Family	Fraction MultiFund
1994	27	7	60	47	17	22	40	176	172	16	584	423	1.38	0.24
1995	34	6	87	67	24	32	55	223	203	24	755	514	1.47	0.29
1996	38	9	110	72	30	38	48	293	180	34	852	600	1.42	0.28
1997	45	14	134	110	43	54	67	393	196	40	1,096	731	1.50	0.33
1998	53	15	157	123	69	60	76	494	208	49	1,304	853	1.53	0.34
1999	58	19	159	136	90	66	76	622	197	69	1,492	958	1.56	0.34
2000	73	23	196	134	103	68	64	760	195	82	1,698	1,066	1.59	0.34
2001	93	16	216	117	136	84	65	815	193	95	1,830	1,118	1.64	0.35
2002	110	14	239	122	155	100	91	860	187	115	1,993	1,190	1.67	0.35
2003	114	17	256	142	164	121	104	930	220	131	2,199	1,253	1.75	0.38
2004	107	12	283	175	178	135	126	987	261	155	2,419	1,314	1.84	0.39
2005	95	11	278	219	198	125	139	1,044	265	174	2,548	1,358	1.88	0.39
2006	76	9	243	237	198	106	137	980	255	226	2,467	1,301	1.90	0.39
2007	68	13	247	321	184	105	138	1,049	276	302	2,703	1,357	1.99	0.40
2008	46	13	180	328	137	91	159	1,007	274	310	2,545	1,282	1.99	0.39
2009	46	9	176	348	128	87	247	927	287	748	3,003	1,271	2.36	0.42
2010	35	4	165	317	115	105	250	807	284	863	2,945	1,195	2.46	0.41
2011	30	5	144	295	114	133	297	722	269	1,198	3,207	1,070	3.00	0.42
2012	26	3	126	218	88	110	256	635	267	1,070	2,799	939	2.98	0.42
2013	28	2	110	170	67	82	204	523	201	933	2,320	782	2.97	0.42

Table 2.2: Inceptions and Flows over time

We report the flows to hedge fund inceptions and existing funds by year. The first column gives the number of fund inceptions in each year. The second column gives the number of new funds in a given year divided by the total number of funds in the universe at the beginning of the year. The third column is the sum of the initial AUMs of the fund inceptions in each year (all AUMs and flows are in billions of US dollars). Initial AUM is defined as the first nonmissing assets under management reported by the fund within 3 months of the fund's inception. Funds with missing AUM are eliminated. The next three columns give the sum of monthly flows in that year to funds that are a year old or younger in each of the year's months, between the age of 1 and 5 years, and over the age of 5 years. Flows are computed monthly fund by fund using

$$\mathrm{Flow}_t = \mathrm{AUM}_t - \mathrm{AUM}_{t-1} \cdot (1 + r_{t-1,t})$$

Total Flows gives the sum of the previous four columns.

Year	New Funds	New /Total	New Fund AUM	Age <=1yr	1yr< Age <=5yr	Age >5yr	Total Flows
1994	208	0.26	1.14	1.75	0.42	0.35	3.66
1995	224	0.22	1.46	1.11	-1.01	-0.99	0.57
1996	303	0.25	1.21	1.92	2.53	-0.73	4.93
1997	317	0.22	2.34	4.85	8.10	-0.49	14.81
1998	308	0.18	1.52	4.44	4.60	-0.97	9.58
1999	389	0.20	3.89	3.96	6.68	-9.04	5.49
2000	388	0.18	3.06	7.50	8.87	-7.82	11.61
2001	444	0.18	3.68	8.92	14.51	3.10	30.21
2002	510	0.18	3.76	9.80	9.03	-3.94	18.64
2003	616	0.18	5.93	11.15	23.41	13.81	54.30
2004	781	0.19	8.35	19.67	31.46	26.35	85.83
2005	824	0.18	7.77	19.25	13.27	-11.93	28.35
2006	771	0.15	7.32	17.56	34.51	15.05	74.44
2007	840	0.16	12.46	19.77	34.19	17.24	83.66
2008	690	0.13	8.35	8.95	-32.58	-48.99	-64.28
2009	678	0.13	8.20	4.92	-15.75	-28.43	-31.06
2010	578	0.11	8.16	8.05	6.55	-8.37	14.38
2011	430	0.08	12.29	4.97	1.48	-7.44	11.30
2012	285	0.06	7.80	3.57	-6.63	-11.25	-6.51
2013	113	0.03	1.68	2.33	-9.94	-11.26	-17.19

Table 2.3: Inceptions in New and Existing Families by Year

We partition the set of all inceptions into inceptions that start new families and those that add a fund to an existing family. The second and third columns report the number of inceptions in each partition by inception year. In the fourth column, we report the number of inceptions in existing families that are not 'clones' of existing funds. That is, the new fund is in a different strategy than all existing funds in the family or it has a return correlation below 90% with all existing funds in the family. Funds with fewer than 12 observations are excluded. The fifth column gives the total number of inceptions in each year, with no requirement of valid AUM at inception.

Year	New Families	Inceptions in Existing Families	Non-Clone Inceptions	Total Inceptions
1994	157	52	32	209
1995	163	68	30	231
1996	217	84	35	301
1997	220	111	53	331
1998	230	89	34	319
1999	285	128	64	413
2000	285	127	65	412
2001	328	189	80	517
2002	368	222	107	590
2003	420	300	121	720
2004	463	415	160	878
2005	487	452	173	939
2006	474	412	196	886
2007	393	592	289	985
2008	344	452	280	796
2009	293	474	258	767
2010	217	459	266	676
2011	189	317	156	506
2012	109	197	88	306
2013	30	90	41	120

Table 2.4: Logistic Regression of Inception Counts on Family/Strategy Variables

This table reports logistic regression results of the dummy variable for whether there was an inception in a given family/quarter/strategy on characteristics of that family, quarter, and strategy category. The logistic regression equation is

$$\text{Inception}_{i,j,t} = \Lambda(\beta \times X_{i,t-1} + \psi \times Y_{k,t-1}) + \varepsilon_{i,j,t}$$

Where $\Lambda(\cdot)$ represents the logistic link function and $X_{i,t}$ is a vector of fund-specific variables for fund i in year t-1 and $Y_{k,t-1}$ is a vector of strategy-specific variables for strategy k (to which fund i belongs) in year t-1. $\text{Inception}_{i,j,t}$ is 1 if there was an inception in strategy i, family j, year t. Otherwise it is zero. Explanatory variables are as follow: Strategy return is the average monthly return to an equal-weighted portfolio of funds in a given strategy in year t-1. Family return is similarly defined. Strategy and family volatility are computed from equal-weighted portfolios of funds from t-2 to t. Strategy and family AUM are the sum of reported assets under management in December of year t-1. Family assets in Same Strategy is the sum of reported assets in strategy i and family j at the end of year t-1. Normalized strategy inceptions is the number of inceptions by strategy i in year t-1, normalized by the number of funds in strategy i at the end of year t-2. Large family open is set to 1 if one of the largest 8 hedge fund families had an inception in strategy i in year t-1. Inceptions in family is the count of inceptions in family j in year t-1. In all models, year-fixed effects are included as yearly dummy variables in the regression. Models 1-4 include only families that existed in the previous year (this investigates the likelihood of inception in existing families) while model 5 adds families that did not exist in the previous year and sets lagged-family variables to 0. In specification (5), a dummy variable set to 1 if the family previously exist is added as well. T statistics are in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Families Existing in the Previous Year				Existing Plus New Families
	(1)	(2)	(3)	(4)	(5)
Strategy Return	29.203***		29.245***	28.531***	36.191***
	(6.62)		(6.59)	(6.44)	(14.67)
Strategy Inceptions	0.788***		0.792***	0.791***	0.892***
	(2.97)		(2.97)	(2.97)	(6.58)
Strategy Volatility	0.449		0.534	0.541	0.897**
	(0.57)		(0.68)	(0.68)	(2.09)
Strategy AUM	0.001***		0.001***	0.001***	0.002***
	(16.47)		(15.91)	(15.93)	(28.68)
Strategy Flow	-4.345		-4.495	-4.423	-11.065***
	(-1.36)		(-1.40)	(-1.38)	(-6.30)
Strategy Large Family Open	0.186**		0.178*	0.175*	0.250***
	(2.01)		(1.91)	(1.87)	(4.53)
Family Return		3.959**	3.864**	3.817**	3.388*
		(2.06)	(2.01)	(1.99)	(1.86)
Family Inceptions		0.175***	0.168***	0.168***	0.162***
		(12.17)	(11.62)	(11.62)	(11.32)
Family Volatility		-1.136***	-1.124***	-1.128***	-1.174***
		(-3.76)	(-3.71)	(-3.72)	(-3.99)
Family AUM		-0.188***	-0.125***	-0.124***	-0.119***
		(-3.97)	(-2.77)	(-2.77)	(-2.67)
Family Assets in Same Strategy		0.529***	0.400***	0.328***	0.326***
		(9.05)	(7.17)	(4.67)	(4.68)
Family Flow		0.321***	0.328***	0.328***	0.344***
		(6.06)	(6.15)	(6.16)	(6.52)
Strategy Return * Family Assets				12.179**	11.022*
				(2.02)	(1.86)
Family Already Existed					-2.179***
					(-37.10)
Fixed Effect: Year	Yes	Yes	Yes	Yes	Yes
Pseudo R-squared	4.90%	3.90%	7.90%	7.90%	18.40%
Obs.	132,340	132,340	132,340	132,340	166,710

Table 2.5: Regression of Inception Holding Period Returns on Fund Characteristics

For each inception, excess returns are computed from the month after the inception (the month in which the inception happens is excluded) to 60 months later. These inception holding period returns are then regressed on strategy category and fund variables. The regression equation is

$$\begin{aligned}\text{Performance}_{i,[1,60]} = \alpha + \beta_1 \text{NormFlow}_{i,[-36,-1]} + \beta_2 \text{Inceptions}_{i,[-12,-1]} \\ + \beta_3 \text{Ret}_{i,[-36,-1]} + \beta_4 \text{Vol}_{i,[-24,-1]} + \beta_5 \text{FirstInFamily}_i \\ + \beta_6 \text{NonClone}_i + \beta_7 \text{InceptionAUM}_i + \beta_8 \text{MissingInception}_i + \varepsilon_i\end{aligned}$$

where Performance_i is the 60-month cumulative return to fund i after its inception. NormFlow_i is the normalized flow to the strategy in the year previous to the inception. Inceptions_i represents the number of inceptions in the strategy category containing fund i divided by the number of funds at the end of the previous period. Ret_i is the lagged equal weighted strategy excess return. Vol_i is the volatility of the strategy containing fund i, computed over the previous 24 months before fund i's inception. FirstInFamily_i is an indicator variable set to one when the fund is the first reported fund in its family. NonClone_i is set to one when a fund is not the first in its family, nor do we consider it a clone fund because it either is in a new strategy for the family or has less than 90% correlation with each previously existing fund in the family. InceptionAUM_i is the first reported AUM for a fund, if the reporting month is within 3 months of its inception date, otherwise it is missing. $\text{MissingInception}_i$ is a dummy variable set to one if there is no reported AUM for fund i within its first three months. T statistics are in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	(1)	(2)	(3)	(4)
Strategy Flow	-9.049***			-9.094***
	(-4.92)			(-4.55)
Strategy Inceptions		-4.866		1.982
		(-1.62)		(0.57)
Strategy Return			-8.148*	-4.411
			(-1.86)	(-0.93)
Strategy Volatility	1.595***	1.567***	1.953***	1.789***
	(3.27)	(3.18)	(3.82)	(3.39)
First Fund in Family	0.267***	0.271***	0.267***	0.265***
	(6.57)	(6.64)	(6.55)	(6.52)
Non-clone Inception in Existing Family	0.289***	0.297***	0.287***	0.285***
	(5.97)	(6.10)	(5.94)	(5.85)
Inception AUM (Bil)	-0.695**	-0.693**	-0.681**	-0.689**
	(-2.03)	(-2.02)	(-1.99)	(-2.02)
Missing Inception AUM	0.016	0.025	0.019	0.014
	(0.46)	(0.71)	(0.54)	(0.38)
Fixed Effect: InceptionYear	Yes	Yes	Yes	Yes
R-Squared	8.60%	8.20%	8.20%	8.60%
Obs.	5,102	5,102	5,102	5,102

Table 2.6: Regression of Inception Holding Period Alpha on Fund Characteristics

For each inception, excess returns are computed from the month after the inception (the month in which the inception happens is excluded) to 60 months later. Alpha for these returns is then computed by regressing fund returns on the Fung Hsieh factors. These inception alphas are then regressed on strategy and fund variables.

$$\begin{aligned}\alpha_i = \delta + \beta_1 \mathrm{NormFlow}_{\mathrm{i},[-36,-1]} + \beta_2 \mathrm{Inceptions}_{\mathrm{i},[-12,-1]} \\ + \beta_3 \mathrm{Ret}_{\mathrm{i},[-36,-1]} + \beta_4 \mathrm{Vol}_{\mathrm{i},[-24,-1]} + \beta_5 \mathrm{FirstInFamily}_{\mathrm{i}} \\ + \beta_6 \mathrm{NonClone}_{\mathrm{i}} + \beta_7 \mathrm{InceptionAUM}_{\mathrm{i}} + \beta_8 \mathrm{MissingInception}_{\mathrm{i}} + \varepsilon_{\mathrm{i}}\end{aligned}$$

where α_i is the alpha coefficient from the Fung Hsieh regression over the first 60-months since inception for fund i. $\mathrm{NormFlow}_i$ is the normalized flow to the strategy over the year previous to inception i. $\mathrm{Inceptions}_i$ represents the number of inceptions in the strategy category containing fund i divided by the number of funds at the end of the previous period. Ret_i is the lagged equal weighted strategy excess return. Vol_i is the volatility of the strategy containing fund i, computed over the previous 24 months. $\mathrm{FirstInFamily}_i$ is an indicator variable set to one when the fund is the first reported fund in its family. $\mathrm{NonClone}_i$ is set to one when a fund is not the first in its family, nor do we consider it a clone fund because it either is in a new strategy for the family or has less than 90% correlation with each previously existing fund in the family. $\mathrm{InceptionAUM}_i$ is the first reported AUM for a fund, if the reporting month is within 3 months of its inception date, otherwise it is missing. $\mathrm{MissingInception}_i$ is a dummy variable set to one if there is no reported AUM for fund i within its first three months. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	(1)	(2)	(3)	(4)
Strategy Flow	-2.199**			-1.321
	(-2.07)			(-1.11)
Strategy Inceptions		-2.522		-0.782
		(-1.55)		(-0.41)
Strategy Return			-8.357***	-6.955**
			(-2.89)	(-2.24)
Strategy Volatility	-1.024***	-1.014***	-0.709**	-0.832**
	(-3.31)	(-3.25)	(-2.23)	(-2.53)
First Fund in Family	0.175***	0.175***	0.173***	0.174***
	(6.82).	(6.80)	(6.73)	(6.77)
Non-clone Inception in Existing Family	0.251***	0.253***	0.251***	0.252***
	(8.36)	(8.42)	(8.36)	(8.38)
Inception AUM (Bil)	-0.191	-0.192	-0.190	-0.190
	(-1.43)	(-1.44)	(-1.43)	(-1.43)
Missing Inception AUM	-0.011	-0.009	-0.012	-0.011
	(-0.52)	(-0.41)	(-0.54)	(-0.50)
Fixed Effect: InceptionYear	Yes	Yes	Yes	Yes
R-Squared	2.80%	2.70%	2.80%	2.80%
Obs.	9,738	9,748	9,748	9,738

Table 2.7: Regression of Inception and Non-Inception Portfolio Returns on Risk Factors

We compare the performance of portfolios of new hedge funds (the inception portfolios) with the performance of portfolios comprising seasoned funds and then compare the performance of inceptions staring during periods of low and high investor demand in the strategy category. Strategy demand is proxied by 36 month returns and normalized flows relative to other categories. Subportfolios are created using a 3-month formation period and then a 60-month holding period. Following Jagadeesh and Titman (1993), for each month, the overall portfolio is formed by equal weighting the subportfolios with holding periods in that month. Thus we average over up to 60 portfolios for each monthly return. For an inception portfolio, the subportfolios consist of funds with inceptions during their formation period. For a non-inception portfolio, subportfolios consist of funds without inceptions in those formation period and which meet the age requirement during the formation period. The alpha for each portfolio is the intercept estimate from the equal-weighted portfolio excess returns on the risk factors

$$r_{p,t} = \alpha_p + \beta_{p,1} MKT_t + \beta_{p,2} SMB_t + \beta_{p,3} YLDCHG_t + \beta_{p,4} BAAMTSY_t \\ + \beta_{p,5} PTFSBD_t + \beta_{p,6} PTFSFX_t + \beta_{p,7} PTFSCOM_t + \varepsilon_{p,t},$$

where $r_{p,t}$ is the excess return on each portfolio in month t. The independent variables are the market excess return (MKT), a size factor (SMB), monthly change in the ten-year Treasury constant maturity yield (YLDCHG), monthly change in the Moody's Baa yield less ten-year Treasury constant maturity yield (BAAMTSY), and three trend-following factors: PFTSBD (bond), PFTSFX (currency), and PFTSCOM (commodity). T statistics are in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Inception and Noninception Portfolios			Inception Portfolios		
	Inceptions	Non-Inceptions	Inceptions minus Noninceptions	Low Strategy Demand Portfolio	High Strategy Demand Portfolio	Low Minus High
alpha	0.362***	0.240***	0.122***	0.483***	0.188	0.295**
	(5.73)	(3.37)	(3.76)	(5.23)	(1.57)	(2.03)
MKT	0.228***	0.278***	-0.050***	0.089***	0.272***	-0.182***
	(15.11)	(16.36)	(-6.46)	(4.07)	(9.49)	(-5.28)
SMB	0.140***	0.156***	-0.016*	0.095***	0.125***	-0.030
	(7.31)	(7.25)	(-1.68)	(3.40)	(3.44)	(-0.69)
YLDCHG	-0.516	-0.592*	0.077	-1.220***	0.045	-1.265*
	(-1.65)	(-1.68)	(0.48)	(-2.67)	(0.08)	(-1.77)
BAAMTSY	-2.122***	-2.774***	0.652***	-3.637***	-2.600***	-1.038
	(-5.32)	(-6.18)	(3.19)	(-6.25)	(-3.43)	(-1.14)
PTFSBD	-0.002	-0.003	0.001	0.009	-0.016**	0.026**
	(-0.45)	(-0.62)	(0.48)	(1.46)	(-2.00)	(2.59)
PTFSFX	0.009**	0.014***	-0.005***	0.009*	0.012*	-0.003
	(2.51)	(3.48)	(-2.75)	(1.76)	(1.76)	(-0.34)
PTFSCOM	0.009*	0.014**	-0.005*	0.007	0.006	0.001
	(1.82)	(2.51)	(-1.96)	(0.99)	(0.67)	(0.07)
R-Squared	63.60%	67.00%	25.10%	27.90%	42.40%	17.40%

Table 2.8: Regression of Inception Portfolio Returns from New Family and Existing Family Inceptions on Risk Factors

We compare the performance of inception portfolios where the new fund is the first reported fund in the family and those that are additions to an existing hedge fund family as well as the clone and non-clone inceptions within existing families. Clones funds are defined as inceptions in families with existing funds of the same strategy category and with 90% or higher correlation with those existing funds. Subportfolios are created using a 3-month formation period and then a 60-month holding period. Following Jagadeesh and Titman (1993), for each month, the overall portfolio is formed by equal weighting the subportfolios with holding periods in that month. Thus we average over up to 60 portfolios for each monthly return. For an inception portfolio, the subportfolios consist of funds with inceptions during their formation period. The alpha for each portfolio is the intercept estimate from the equal-weighted portfolio excess returns on hedge fund risk factors

$$r_{p,t} = \alpha_p + \beta_{p,1} MKT_t + \beta_{p,2} SMB_t + \beta_{p,3} YLDCHG_t + \beta_{p,4} BAAMTSY_t$$
$$+ \beta_{p,5} PTFSBD_t + \beta_{p,6} PTFSFX_t + \beta_{p,7} PTFSCOM_t + \varepsilon_{p,t},$$

where $r_{p,t}$ is the excess return on each portfolio in month t. The independent variables are the market excess return (MKT), a size factor (SMB), monthly change in the ten-year Treasury constant maturity yield (YLDCHG), monthly change in the Moody's Baa yield less ten-year Treasury constant maturity yield (BAAMTSY), and three trend-following factors: PFTSBD (bond), PFTSFX (currency), and PFTSCOM (commodity). T statistics are in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	New Family and Existing Family Inceptions			Existing Family Inceptions		
	New Family Inceptions	Existing Family Inceptions	New minus Existing	Non-Clone	Clone	Non-Clone minus Clone
alpha	0.372***	0.338***	0.035	0.408***	0.279***	0.129***
	(6.06)	(4.55)	(0.96)	(5.36)	(3.56)	(2.61)
MKT	0.240***	0.228***	0.012	0.231***	0.233***	-0.003
	(16.36)	(12.88)	(1.38)	(12.70)	(12.46)	(-0.21)
SMB	0.144***	0.135***	0.009	0.143***	0.136***	0.006
	(7.74)	(6.01)	(0.83)	(6.20)	(5.74)	(0.43)
YLDCHG	-0.308	-0.887**	0.579***	-0.979***	-0.760*	-0.219
	(-1.01)	(-2.42)	(3.25)	(-2.60)	(-1.96)	(-0.90)
BAAMTSY	-2.107***	-2.232***	0.125	-1.854***	-2.528***	0.674**
	(-5.43)	(-4.77)	(0.55)	(-3.86)	(-5.11)	(2.16)
PTFSBD	-0.004	-0.001	-0.003	-0.001	-0.001	0.001
	(-0.93)	(-0.12)	(-1.34)	(-0.10)	(-0.21)	(0.18)
PTFSFX	0.009***	0.011***	-0.002	0.012***	0.010**	0.002
	(2.67)	(2.61)	(-0.82)	(2.78)	(2.33)	(0.59)
PTFSCOM	0.006	0.016***	-0.010***	0.013**	0.016***	-0.004
	(1.34)	(2.85)	(-3.58)	(2.21)	(2.78)	(-1.00)
R-Squared	67.20%	55.40%	17.20%	53.70%	54.70%	5.10%

Table 2.9: Strategy Demand and Inception Type Performance Matrix—Risk Adjusted Alpha

We form inception portfolios based on the type of inception and our measure of strategy category demand at the time of inception. Investor demand is proxied by relatively high strategy returns and strategy flows over the previous 36 months. We say that a fund in a strategy in the bottom four deciles by both measures came into being during a period of relatively low investor demand. A fund in the top four deciles came into being during a period of relatively high investor demand. Inception types are Non-Clone, meaning the fund was not the first in its family but it is not closely related to any existing funds in the. New Family funds are the first funds in their reported family. Clone funds are those that are not the first fund in their family and are closely related to existing funds in the family (same strategy and 90% correlation or higher with an existing fund in the family). Subportfolios are created using a 3-month formation period and then a 60-month holding period. Following Jegadeesh and Titman (1993), for each month, the overall portfolio is formed by equal weighting the subportfolios with holding periods in that month. Thus we average over up to 60 portfolios for each monthly return. For an inception portfolio, the subportfolios consist of funds with inceptions during their formation period. The risk adjusted alpha for each portfolio is the intercept estimate from the equal-weighted portfolio excess returns on hedge fund risk factors. The regression equation is

$$r_{p,t} = \alpha_p + \beta_{p,1} MKT_t + \beta_{p,2} SMB_t + \beta_{p,3} YLDCHG_t + \beta_{p,4} BAAMTSY_t$$
$$+ \beta_{p,5} PTFSBD_t + \beta_{p,6} PTFSFX_t + \beta_{p,7} PTFSCOM_t + \varepsilon_{p,t},$$

where $r_{p,t}$ is the excess return on each portfolio in month t. The independent variables are the market excess return (MKT), a size factor (SMB), monthly change in the ten-year Treasury constant maturity yield (YLDCHG), monthly change in the Moody's Baa yield less ten-year Treasury constant maturity yield (BAAMTSY), and three trend-following factors: PFTSBD (bond), PFTSFX (currency), and PFTSCOM (commodity). This table reports the estimated α_p for each portfolio. T statistics are in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Non-Clone	New Family	Clone	Non-Clone minus Clone	New Family minus Clone
Low Strategy Demand	0.469%	0.463%	0.388%	0.005%	0.010%
	(3.751)	(4.899)	(2.033)	(0.044)	(0.062)
High Strategy Demand	0.344%	0.244%	0.199%	0.182%	0.089%
	(2.472)	(2.070)	(1.309)	(1.649)	(0.918)
Low - High	0.193%	0.292%	0.370%	0.375%	0.380%
	(1.096)	(2.064)	(1.619)	(2.115)	(2.401)

Table 2.10: Strategy Demand and Inception Type Performance Matrix—Risk Coefficients

We form inception portfolios based on the type of inception and our measure of strategy category demand at the time of inception. Investor demand is proxied by relatively high strategy returns and flows during the previous 36 months. We say that a fund in the bottom four deciles of strategies came into being during a period of relatively low investor demand. A fund in the top four deciles of strategies came into being during a period of relatively high investor demand. Inception types are Non-Clone, meaning the fund was not the first in its family but it is not closely related to any existing funds in the family. New Family funds are the first funds in their reported family. Clone funds are those that are not the first fund in their family and are closely related to existing funds in the family (same strategy and 90% correlation or higher with an existing fund in the family). Subportfolios are created using a 3-month formation period and then a 60-month holding period. Following Jagadeesh and Titman (1993), for each month, the overall portfolio is formed by equal weighting the subportfolios with holding periods in that month. For an inception portfolio, the subportfolios consist of funds with inceptions during their formation period. The regression equation is

$$r_{p,t} = \alpha_p + \beta_{p,1} MKT_t + \beta_{p,2} SMB_t + \beta_{p,3} YLDCHG_t + \beta_{p,4} BAAMTSY_t + \beta_{p,5} PTFSBD_t + \beta_{p,6} PTFSFX_t + \beta_{p,7} PTFSCOM_t + \varepsilon_{p,t},$$

where $r_{p,t}$ is the excess return on each portfolio in month t. The independent variables are the market excess return (MKT), a size factor (SMB), monthly change in the ten-year Treasury constant maturity yield (YLDCHG), monthly change in the Moody's Baa yield less ten-year Treasury constant maturity yield (BAAMTSY), and three trend-following factors: PFTSBD (bond), PFTSFX (currency), and PFTSCOM (commodity). This table reports the regression coefficients for each portfolio. T statistics are in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Low NonClone Minus High NonClone	Low New minus High New	Low Clone Minus High Clone	Low NonClone Minus High Clone	Low New Minus High Clone
alpha	0.193	0.292**	0.370	0.375**	0.380**
	(1.10)	(2.06)	(1.62)	(2.12)	(2.40)
MKT	-0.136***	-0.204***	-0.134**	-0.086**	-0.169***
	(-3.25)	(-6.06)	(-2.46)	(-2.04)	(-4.48)
SMB	-0.081	-0.021	-0.176**	-0.142***	-0.069
	(-1.52)	(-0.50)	(-2.55)	(-2.64)	(-1.45)
YLDCHG	-3.110***	-0.594	-2.692**	-2.827***	-0.560
	(-3.57)	(-0.85)	(-2.38)	(-3.22)	(-0.71)
BAAMTSY	-2.213**	-0.091	-2.529*	-2.291**	-0.648
	(-1.99)	(-0.10)	(-1.75)	(-2.05)	(-0.65)
PTFSBD	0.039***	0.023**	0.044***	0.049***	0.026**
	(3.22)	(2.32)	(2.79)	(4.03)	(2.40)
PTFSFX	-0.002	-0.002	0.016	0.012	0.006
	(-0.20)	(-0.30)	(1.20)	(1.21)	(0.61)
PTFSCOM	0.013	0.001	0.017	-0.003	-0.014
	(0.95)	(0.06)	(0.98)	(-0.19)	(-1.19)
R-Squared	17.50%	20.00%	15.40%	18.70%	14.20%

Table 2.11: High and Low investor Demand Inception Portfolios (No Backfill)

We compare the performance of inception portfolios for funds that came at times of low investor demand and those that came at time of high investor demand. Investor demand is proxied by strategy inceptions over the previous 12 months. We say that an inception in the bottom four deciles of strategies occurred during a period of relatively low investor demand. Likewise we say that an inception in the top four deciles of strategies happened during a period of relatively high investor demand. Subportfolios are created using a 3-month formation period and then a 60-month holding period. Following Jagadeesh and Titman (1993), for each month, the overall portfolio is formed by equal weighting the subportfolios with holding periods in that month. Thus we average over up to 60 portfolios for each monthly return. For an inception portfolio, the subportfolios consist of funds with inceptions during their formation period. We do not include any backfilled returns (returns previous to the date a fund started reporting to TASS). The alpha for each portfolio is the intercept estimate from the equal-weighted portfolio excess returns on hedge fund risk factors

$$r_{p,t} = \alpha_p + \beta_{p,1} MKT_t + \beta_{p,2} SMB_t + \beta_{p,3} YLDCHG_t + \beta_{p,4} BAAMTSY_t$$
$$+ \beta_{p,5} PTFSBD_t + \beta_{p,6} PTFSFX_t + \beta_{p,7} PTFSCOM_t + \varepsilon_{p,t},$$

where $r_{p,t}$ is the excess return on each portfolio in month t. The independent variables are the market excess return (MKT), a size factor (SMB), monthly change in the ten-year Treasury constant maturity yield (YLDCHG), monthly change in the Moody's Baa yield less ten-year Treasury constant maturity yield (BAAMTSY), and three trend-following factors: PFTSBD (bond), PFTSFX (currency), and PFTSCOM (commodity). T statistics are in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Low Demand Portfolio	High Demand Portfolio	Spread
alpha	0.489***	0.196	0.293**
	(5.26)	(1.63)	(2.00)
MKT	0.087***	0.274***	-0.187***
	(3.91)	(9.53)	(-5.37)
SMB	0.096***	0.126***	-0.030
	(3.40)	(3.46)	(-0.69)
YLDCHG	-1.248***	0.024	-1.273*
	(-2.72)	(0.04)	(-1.76)
BAAMTSY	-3.654***	-2.608***	-1.046
	(-6.24)	(-3.43)	(-1.13)
PTFSBD	0.011*	-0.016**	0.027***
	(1.66)	(-1.99)	(2.69)
PTFSFX	0.010*	0.012*	-0.002
	(1.82)	(1.68)	(-0.23)
PTFSCOM	0.007	0.007	0.000
	(0.94)	(0.72)	(0.00)
R-Squared	27.50%	42.60%	18.00%

Table 2.12: High and Low investor Demand Inception Portfolios (Alternate formation)

We compare the performance of inception portfolios for funds that came at times of low investor demand and those that came at time of high investor demand. Investor demand at the strategy category level is proxied by two variables: relatively high strategy normalized flows (relative to other strategies) over the previous 36 months and relatively strategy returns over the previous 36 months. We say that an inception in the bottom four deciles of strategies in both respects happened during a period of relatively low investor demand. A fund in the top four deciles of strategies in both respects came into being during a period of relatively high investor demand. Subportfolios are created using a 3-month formation period and then a 60-month holding period. Following Jagadeesh and Titman (1993), for each month, the overall portfolio is formed by equal weighting the subportfolios with holding periods in that month. Thus we average over up to 60 portfolios for each monthly return. For an inception portfolio, the subportfolios consist of funds with inceptions during their formation period. The alpha for each portfolio is the intercept estimate from the equal-weighted portfolio excess returns on hedge fund risk factors

$$r_{p,t} = \alpha_p + \beta_{p,1} MKT_t + \beta_{p,2} SMB_t + \beta_{p,3} YLDCHG_t + \beta_{p,4} BAAMTSY_t$$
$$+ \beta_{p,5} PTFSBD_t + \beta_{p,6} PTFSFX_t + \beta_{p,7} PTFSCOM_t + \varepsilon_{p,t},$$

where $r_{p,t}$ is the excess return on each portfolio in month t. The independent variables are the market excess return (MKT), a size factor (SMB), monthly change in the ten-year Treasury constant maturity yield (YLDCHG), monthly change in the Moody's Baa yield less ten-year Treasury constant maturity yield (BAAMTSY), and three trend-following factors: PFTSBD (bond), PFTSFX (currency), and PFTSCOM (commodity). T statistics are in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Low Demand Portfolio	High Demand Portfolio	Spread
alpha	0.453***	0.300***	0.154**
	(7.08)	(3.71)	(2.03)
MKT	0.159***	0.233***	-0.074***
	(10.39)	(12.11)	(-4.11)
SMB	0.089***	0.155***	-0.066***
	(4.58)	(6.35)	(-2.88)
YLDCHG	-0.852***	-0.117	-0.735*
	(-2.69)	(-0.29)	(-1.96)
BAAMTSY	-2.898***	-1.929***	-0.969**
	(-7.17)	(-3.79)	(-2.02)
PTFSBD	-0.004	-0.007	0.003
	(-0.86)	(-1.19)	(0.53)
PTFSFX	0.012***	0.007	0.006
	(3.38)	(1.41)	(1.35)
PTFSCOM	0.006	0.006	0.000
	(1.28)	(1.03)	(-0.01)
R-Squared	52.20%	54.40%	13.30%

Table 2.13: Clone and Non Clone Inception Portfolios (Alternate formation)

We compare the performance of inception portfolios in existing families where the new fund is considered a clone of an existing fund in the family and those that are considered non-clones. Clones funds are defined as inceptions in families with existing funds of the same strategy category and with 85% or higher correlation with those existing funds. Subportfolios are created using a 3-month formation period and then a 60-month holding period. Following Jagadeesh and Titman (1993), for each month, the overall portfolio is formed by equal weighting the subportfolios with holding periods in that month. Thus we average over up to 60 portfolios for each monthly return. For an inception portfolio, the subportfolios consist of funds with inceptions during their formation period. The alpha for each portfolio is the intercept estimate from the equal-weighted portfolio excess returns on hedge fund risk factors

$$r_{p,t} = \alpha_p + \beta_{p,1} MKT_t + \beta_{p,2} SMB_t + \beta_{p,3} YLDCHG_t + \beta_{p,4} BAAMTSY_t + \beta_{p,5} PTFSBD_t + \beta_{p,6} PTFSFX_t + \beta_{p,7} PTFSCOM_t + \varepsilon_{p,t},$$

where $r_{p,t}$ is the excess return on each portfolio in month t. The independent variables are the market excess return (MKT), a size factor (SMB), monthly change in the ten-year Treasury constant maturity yield (YLDCHG), monthly change in the Moody's Baa yield less ten-year Treasury constant maturity yield (BAAMTSY), and three trend-following factors: PFTSBD (bond), PFTSFX (currency), and PFTSCOM (commodity). T statistics are in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Non-Clone	Clone	Non-Clone minus Clone
alpha	0.407***	0.302***	0.106*
	(5.23)	(3.72)	(1.97)
MKT	0.227***	0.236***	-0.009
	(12.26)	(12.23)	(-0.68)
SMB	0.131***	0.139***	-0.007
	(5.61)	(5.69)	(-0.46)
YLDCHG	-0.962**	-0.814**	-0.148
	(-2.51)	(-2.04)	(-0.56)
BAAMTSY	-1.834***	-2.537***	0.703**
	(-3.76)	(-4.99)	(2.09)
PTFSBD	-0.002	-0.001	-0.001
	(-0.37)	(-0.15)	(-0.30)
PTFSFX	0.012***	0.012**	0.001
	(2.75)	(2.49)	(0.23)
PTFSCOM	0.013**	0.018***	-0.005
	(2.21)	(2.92)	(-1.20)
R-Squared	52.30%	54.50%	4.60%

Chapter 3: Return Smoothing, Liquidity Costs, and Investor Flows: Evidence from a Separate Account Platform

Abstract

We use a new dataset of hedge fund returns from a separate account platform to examine (1) how much of hedge fund return smoothing is due to main-fund specific factors, such as managerial reporting discretion (2) the costs of removing hedge fund share restrictions. These accounts trade *pari passu* with matching hedge funds but feature third-party reporting and permissive share restrictions. We use these properties to estimate that 33% of reported smoothing is due to managerial reporting methods. The platform's fund-level liquidity is associated with costs of 1.7% performance reduction on an annual basis. Investor flows chase monthly past performance on the platform but not in the associated funds.

Keywords: hedge funds, separate accounts, return smoothing, share restrictions
JEL classifications: G23, G11

3.1 Introduction

Hedge funds returns are conspicuous among investment vehicles because their voluntarily reported (typically non-audited) returns consistently manifest positive serial correlation. Share restrictions, such as lockup periods and redemption notice periods, and delayed reporting may prevent investors from arbitraging this serial correlation away, but what is its cause? Is serial correlation inherent to the underlying assets hedge funds hold or are managers smoothing returns before reporting them to data providers? Stringent share restrictions may have obvious costs to investors, but how significant are their benefits? If fund managers were to allow frequent subscriptions and redemptions at low cost, would hedge fund investor wealth significantly suffer?

Fundamental and important as these questions are, the standard commercial hedge fund databases have a difficult time answering them. We observe serial correlation in reported hedge fund returns, but typically do not observe the trades or mark-to-market methods, so we typically cannot know how a third party would report monthly returns for the same funds. In essence, we do not know whether serial correlation is present in a fund's assets or only in the reported returns. Similarly, we observe varying share restrictions across funds, but there are many other significant differences between funds that obfuscate the causative relation, if any, between share restrictions and fund performance.

In this paper we use the unique properties of the Lyxor separate account platform to directly examine why reported hedge fund returns are serially correlated or smoothed. The Lyxor platform manages hedge fund separate accounts that are linked to the main hedge funds by agreement with the fund managers that they will make the same trades in the separate accounts as in the main fund. Thus we say that separate accounts are traded *pari passu* with the associated main funds. Although separate accounts and main funds trade the same assets, separate account returns are computed by a third party associated with the platform and reported to investors at a relatively high frequency. This removes the opportunity for hedge fund managers to manipulate or misreport the returns of the separate account. Since the underlying assets of a separate account match those of the associated main fund, any smoothing found in the returns of the separate account can be attributed to the underlying assets or hedge fund strategy.

The properties of the separate account platform also allow us to directly observe the performance implications of share restrictions. Separate accounts have very different share restrictions and effective fund-level liquidity from the main funds but have the same underlying assets, strategy, and manager[1]. As a result, we can attribute differences in performance between the main funds and associated separate accounts to the costs associated with the separate accounts' fund flows, which are largely absent in the main funds because of those funds' share restrictions. Any performance difference between the funds and their matching accounts is evidence of the effect of share restrictions on investor wealth.

This paper uses a new and more direct approach examining the effect of managerial reporting on return smoothing. Hedge fund returns are typically calculated and voluntarily reported by the funds themselves and there is evidence that fund managers can manipulate the calculation of these returns to improve apparent fund performance or to maximize the fees they collect (e.g., see Agarwal, Daniel, and Naik, 2011; Bollen and Pool, 2009). On the other hand, there could be less pernicious explanations for this serial correlation, such as the funds inheriting serially correlated returns from underlying assets that manifest this property as in Getmansky, Lo, and Makarov (2004). By comparing the performance of the hedge fund portfolio as self-reported by the manager and as reported by the separate account platform, we can quantify the effect of managerial discretion on reported serial correlation in a more direct way than previous research has been able to do. Since the main fund and associated separate account have the same manager, assets, and trades, the variables that prevent researchers from cleanly disentangling asset-induced and managerial smoothing are eliminated in this context. We can then decompose the smoothing in main fund returns (which is what is normally observable) into the portion attributable to the underlying assets/strategy and the portion that may be attributed to managerial reporting choices.

We discuss the mechanisms that may cause the gap in performance and serial correlation between a hedge fund and its associated separate account, including differences in reporting practices and valuation methods, temporary differences in portfolio holdings due to flow-induced trading, cash management, and serial correlation due to price pressure caused by flow-induced trading.

This paper is also unique in that it directly investigates the relation between share restrictions

[1]From the point of view of an investor, cash invested in a hedge fund may be impossible to liquidate rapidly due to share restrictions even if the fund itself holds very liquid assets. We distinguish the two concepts by saying that such an investment has low *fund-level* liquidity while the assets being held by the fund have high *asset-level* liquidity. Fund-level liquidity is determined by the provisions agreed to by fund management and investors, not by the characteristics of the underlying assets.

and fund performance using a matched sample in which associated funds differ significantly in their share restrictions but not in their investment choices. Liang (1999) and Aragon (2007) investigate the effect of share restrictions and find that funds with low fund-level liquidity by virtue of these restrictions outperform their more liquid peers. The later study finds evidence that funds with less fund-level liquidity tend to hold less liquid assets and suggests that share restrictions provide managers the freedom to undertake profitable long-horizon strategies that may not be amenable to short-term liquidation. Joenvaara and Kosowski (2013) examine the performance difference between UCITS and non-UCITS hedge funds (the former often have fewer share restrictions) and find that non-UCITS funds typically outperform and that the difference is explained by differences in liquidity and share restrictions[2]. They group funds by characteristics such as share restrictions, domicile, and geography for comparison. In contrast to previous studies we are able to directly compare matching funds, with the same manager and strategy, in two share restriction regimes: the main fund and the platform separate account.

Finally, we examine the effect of share restrictions on investor flows: Investors on the separate account platform can move money from one fund to another within the platform as frequently as once a week, whereas investors in the main hedge funds are subject to the redemption frequency requirements (the average redemption frequency is about two months for our sample funds) and other constraints imposed by the main fund. The lack of these barriers on the separate account platform drastically increases investors' ability to chase hedge fund performance.

Methodologically, we construct a merged main fund and separate account return dataset in which we match contemporaneous returns from the separate accounts with their associated main funds. We examine the difference between returns as reported by hedge fund managers and those realized by separate accounts to determine the degree to which managerial discretion in valuing assets plays a role in the smoothing of main funds' reported returns as well as the effect of the share restrictions on the main fund performance.

Estimating the moving average process of Lyxor and main funds, we find that matched main funds have an average first-order moving average coefficient of 0.182 while the associated Lyxor accounts have, on average, a statistically different coefficient of 0.121. Since underlying asset

[2] Undertakings for Collective Investment in Transferable Securities (UCITS) funds conform to a set of European Union directives that allow funds to operate freely in any EU state on the basis of authorization from a single state.

illiquidity should be the same for a fund and its separate account, we conclude that 33% of reported hedge fund return smoothing is due to managerial discretion in return reporting and 67% of reported return smoothing is due to the properties of the underlying assets and other factors common to the main fund and separate account. The proportion of moving average smoothing due to managerial discretion is higher for funds with greater barriers to liquidity as measured by the total time to liquidation caused by the fund's share restrictions.

We also find that Lyxor separate accounts underperform their matched main funds over our sample by 1.7% annually, suggesting that the increased fund-level liquidity of the separate accounts is quite costly. The magnitude of this difference is also positively associated with the higher share restrictions. In other words the imposition of share restrictions, while it reduces investor flexibility, also shields investors from significant fund-liquidity costs.

Examining the relation between fund performance and subsequent flows, we find strong evidence that investors chase fund performance on the Lyxor platform at the one, two, and three-month horizons, while evidence among the main funds is weaker and restricted to the two-month horizon. This result suggests that some investors in the main funds would move cash to chase high-performance if they were not prevented from doing so by the funds' share restrictions.

The rest of the paper proceeds as follows: Section 3.2 develops our hypotheses of interest and place them in the context of the existing literature. Section 3.3 describes the Lyxor separate account platform, our dataset, and the methodology used to construct a sample appropriate for testing. Section 3.4 then outlines the econometric tests used and discusses the results. Section 3.5 concludes the paper.

3.2 Hypothesis Development

3.2.1 Third-Party Valuation, Fund Liquidity, and Return Smoothing

Self-reported hedge fund returns are serially correlated (see Asness, Krail, and Liew, 2001; Getmansky, Lo, and Makarov, 2004; Loudon, Okunev, and White, 2006). The fact that these returns are voluntarily reported and that independent valuation of the underlying securities in the portfolio are not available has led some academics to suggest that hedge fund managers purposely manipulate reported net asset values (NAV). For example, Bollen and Pool (2008) discusses how hedge fund managers have incentive to report positive returns right away while inducing serial correlation by delaying negative fund returns. Moreover, Bollen and Pool (2009) find evidence in the pooled distribution of hedge fund returns that the reported small gains far exceed the reported small losses and this discontinuity in returns is absent in the three months preceding an audit. Further, the discontinuity disappears when using bimonthly returns, indicating that the reported returns are subsequently reversed. Agarwal, Daniel, and Naik (2011) show abnormally positive hedge fund returns in December, a month in which many fee schedules crystallize, consistent with earnings manipulation for the purpose of maximizing fees.

Manipulation of hedge fund reported returns is possible because hedge funds have a great deal of discretion in the information they reveal and to whom they reveal it. The performance information in the main hedge fund databases is voluntarily reported by hedge fund management after the fact. Presumably, the manager would avoid reporting returns that are very different in the long run from what was actually achieved because some investors in the fund may have access to the databases and notice the discrepancy, particularly if they liquidated or added cash to the fund and are aware of the NAV used. However, if there are months in which no assets are added to or removed from the fund, there would be no way for an investor to know whether the reported returns have been manipulated nor how up to date the NAV really is. These characteristics of self-reported returns are common among hedge funds, but not possible among the separate accounts on the Lyxor platform. For one thing, investors on the platform have up to weekly liquidity opportunities and weekly performance reporting, so any manipulation of reported returns would also imply liquidations occurring at the manipulated NAV, a situation that would lead to wealth transfers between fund investors. More importantly, the NAV calculation on the platform is performed by a third party without the incentive to alter returns. Therefore Lyxor account returns are free of managerial discretion in reporting.

There are alternative explanations for serial correlation that do not involve deliberate manipulation of returns by fund managers. Getmansky, Lo, and Makarov (2004) present an econometric model of serial correlations and illiquidity and suggest that smoothed (and predictable) hedge fund returns are largely a result of illiquid underlying securities. The stale pricing of these investments leads mechanically to apparently smoothed fund returns in the same way that non-synchronous trading can cause spurious serial correlation in equity indexes. They highlight linear extrapolation of prices of thinly traded securities, smoothed broker-dealer quotes, and trading restrictions as mechanisms for hedge fund asset smoothing. These mechanisms should affect return calculation for the separate accounts and associated main funds equally if they are using the same methods to compute returns. Alternatively, if fund managers choose a valuation method that induces greater serial correlation in reported returns than what Lyxor, as a third party, chooses or if they directly manipulate return numbers, we would expect to see a difference in serial correlation between main funds and separate accounts. These issues lead to the following hypothesis.

Hypothesis 3.1 *To the extent that hedge fund managers use discretion to game evaluation by smoothing returns, we will observe greater smoothing in main fund returns than in corresponding Lyxor returns.*

Because the assets and trading in the Lyxor accounts match those of the corresponding main funds, we believe the asset-induced smoothing due to stale prices and non-synchronous trading to be approximately the same for the main funds as for the Lyxor separate accounts. Therefore we can decompose the smoothing in the main funds into two components: the portion inherited from the underlying assets (also the amount we find in the Lyxor account) and the portion induced by manager valuation (the difference in smoothing between the main funds and Lyxor accounts).

If investors are primarily concerned with accurate reporting when making subscriptions or redemptions, then managers with greater share restrictions will have greater managerial discretion in reporting returns. This leads to the following hypothesis.

Hypothesis 3.2 *If share restrictions facilitate managerial discretion, then funds with more stringent share restrictions will have greater smoothing than those with less.*

3.2.2 Hedge Fund Liquidity and Performance

In the context of this paper, hedge fund liquidity refers to the difficulty or ease with which a hedge fund's investors can make or redeem an investment in the fund. Since hedge fund liquidity is defined by the share restrictions agreed upon by the hedge fund manager and investors, it is related to the liquidity of the underlying assets only through the views of the fund management, since they decide what terms the fund will offer investors. Fund liquidity should not be confused with the liquidity of the fund's underlying assets, which is defined by costs faced by the hedge fund managers attempting to enter into or liquidate positions in underlying assets in order to meet subscription/redemption, rebalancing, or other trading demands.

One important research question is whether these liquidity restrictions are justified and beneficial to investors. Aragon (2007) compares the returns of hedge funds with and without lockup periods and other restrictions and finds that after accounting for illiquidity provisions, hedge fund returns no longer have a positive alpha. This result supports the notion that hedge funds are, effectively, funds for long-horizon investors who are willing to trade liquidity for higher returns. On the other hand, Ang and Bollen (2010) showed that share restrictions such as lockups and notice periods can be very costly to investors when fund-level liquidity is viewed as a real option.

Since Lyxor accounts and their corresponding main funds differ primarily in their fund-level liquidity restrictions (with Lyxor accounts being universally more liquid), we can test hypotheses about the difference between Lyxor and main fund returns by comparing matching Lyxor and main fund returns. Hence, we develop the following testable hypothesis.

Hypothesis 3.3 *If greater hedge fund liquidity imposes higher costs on the fund, Lyxor accounts will underperform their corresponding main funds.*

Further, if liquidity restrictions are being rationally imposed by fund managers in response to the expected costs of flows in their funds, we expect the performance difference to be related to the main fund's liquidity restrictions.

Hypothesis 3.4 *If share restrictions are put in place in the main fund in order to mitigate costs associated with investor flows, then funds with greater share restrictions in the main fund will show a greater performance difference between the main fund and separate account.*

3.2.3 Fund Liquidity and the Return-Flow Relation

Flows of new money from investors in the mutual fund industry tend to chase past performance (see Sirri and Tufano, 1998). Zheng (1999) shows that this phenomenon may be rational, since mutual funds with good performance tend to continue to have good performance in the short run—a finding similar to the "hot hands" effect of Hendricks, Patel, and Zeckhauser (1993).

In empirical tests of hot hands performance persistence among hedge funds, Jagannathan, Malakhov, and Novikov (2010) find that 25% of three-year abnormal performance spills over into the following three-year interval. Their investigation looks at the three-year horizon because many funds have lockup periods of two to three years. Agarwal and Naik (2000) examine performance persistence at various frequencies and find quarterly performance persistence to be most robust, with weaker effects at longer horizons.

Investors may be inclined to move cash to take advantage of this performance persistence. Agarwal, Daniel, and Naik (2004) find evidence that at the annual frequency, hedge fund flows follow performance. Investors seeking to chase performance may be impeded by the presence of share restrictions. Aragon, Liang, and Park (2013) examine onshore and offshore hedge funds, the latter of which often have lower share restrictions, and find that flows chase performance more in offshore funds than in their onshore counterparts. This difference may be due to the effect of share restrictions.

It is generally difficult for investors to move cash into and out of hedge funds quickly enough to chase performance at the monthly frequency since many funds have quarterly redemption dates, quarterly reporting, or delayed reporting. Other liquidity provisions can make even longer-horizon performance chasing impossible among some hedge funds. In contrast, once cash is on the Lyxor platform, it can be moved from one account to another as often as weekly and without transaction costs. These features allow Lyxor customers to chase performance to a much greater degree than investors in the main funds can, leading to the following hypothesis.

Hypothesis 3.5 *If hedge fund investors seek to chase fund performance, we expect a stronger positive relation between fund performance and subsequent flows for Lyxor funds than for main funds, which provide fewer opportunities for flows.*

3.3 Data Description

3.3.1 The Lyxor Separate Account Platform

We utilize a unique dataset taken from separate accounts on the Lyxor separate account platform, a wholly owned subsidiary of Société Générale. At the end of 2011, Lyxor was the largest and best-known separate account platform in the world, with over 100 separate accounts available for investment and $10 billion in assets under management (AUM). Figure 3.1 shows the growth in Lyxor platform AUM over time. Aside from Lyxor, there are several other large and growing separate account platforms managed by Deutche Bank, Man Group, AlphaMetrix, Goldman Sachs, UBS, and other institutions. While some institutional details vary from one separate account platform to another, they each provide the same general features: third party validation/evaluation and increased ability to move cash from one fund to another quickly and at low cost. Total assets across all separate account platforms is estimated to be close to $100 billion.

Lyxor's separate account platform contracts with operating hedge funds (the main funds) to open separate accounts that receive funds from investors on the Lyxor platform. These funds are then traded *pari passu* with the assets in the main fund[3]. Investors in the platform are subject to the same legal requirements as those of a hedge fund: They must be legally accredited investors.

The Lyxor separate account platform was created in 1998 to provide sophisticated investors access to hedge funds in a manner that (1) allows for diversification across hedge funds without having to meet the minimum investment requirement for each hedge fund, (2) allows greater fund-level liquidity than the main hedge funds, (3) provides investors greater transparency in return reporting, (4) mitigates the risk of fund-level fraud, and (5) increases standardization of the fee structure and account terms across funds. Typical investors are private banks, pension funds, and other institutional investors who have a preference for the liquidity, transparency, or institutional risk characteristics of the platform.

Although the separate account platform has a minimum investment size of $100,000, there is no minimum investment for an allocation to any particular account on the platform. After meeting the overall minimum, an investor may diversify across many separate accounts to create a portfolio with returns similar to a diversified hedge fund index or fund of hedge funds. The separate accounts on the platform charge management and incentive fees, just as the main funds do. Lyxor also has a management fee proportional to the assets invested on the platform, but this fee does not enter into our calculations because our returns are gross of Lyxor fees and net of the fund managers' management and incentive fees. Because Lyxor's fees are not included in this data, the difference between main fund and separate account performance we report is a lower bound on what investors would face.

While main funds often provide self-reported performance information relatively infrequently

[3] Although separate accounts participate in the same trades as the main funds, the returns will not perfectly match those of the main fund because of the timing of flows, differing transaction costs, etc.

(either to investors or to the data providers used by researchers), the Lyxor platform calculates and publishes the NAV and returns of each separate account weekly. Since Lyxor calculates the asset valuations independently, many of the issues associated with returns that are self-reported by the fund manager (e.g., intentional return smoothing) can be avoided. Lyxor monitors the trading and returns of the separate accounts and compares them against those of the main fund to ensure that the funds do not experience style drift and that the accounts trade in step with the main funds.

Cash may be moved from one separate account on the platform to another as often as weekly, after the close of trading each Tuesday, at the published NAV without incurring transaction costs. During our sample period, orders for these transactions had to be submitted with two days' notice—that is, the Friday before the Tuesday of the transaction. This potentially allows investors to move funds quickly from one separate account to another, with performance information from all but the most recent few days. Because the published NAV for each Tuesday is used to compute the value of the shares created and redeemed by platform participants, it need not be considered indicative, as would a self-reported NAV in a hedge fund at the end of a month in which redemptions and subscriptions are not permitted.

As much as possible, subscriptions and redemptions from different platform investors are netted out to avoid forcing funds to liquidate or enter into positions unnecessarily. A handful of funds consistently allow less frequent subscriptions and redemptions, but other funds may occasionally be flagged as having "limited" or "very limited" liquidity. Investors seeking to take large positions in funds flagged in this manner may not get their orders completely filled. Cash not allocated to a separate account by the investor or allocated to an account but not yet invested by the separate account manager is managed by Lyxor and earns a risk-free rate.

The rapid growth in assets managed by separate account platforms highlights the need some investors have to obtain hedge fund returns without being subject to the share restrictions, difficulty of diversification, and perceived institutional risk they would face if they invested in conventional hedge funds. In fact, some institutional investors are prohibited from direct hedge fund investment but are permitted to invest on the platform because the assets are held by Société Générale and evaluated using a third party. In other words Lyxor separate accounts provide a degree of credibility that direct hedge fund investments cannot easily match.

The desirable liquidity and transparency features of the separate account platform do come with some costs to investors. Lyxor's fee, which is not included in the performance we report here, will reduce the net return to investors. Moreover, inasmuch as the managers are not able to cheaply implement the trades necessary to provide investor liquidity, they are instructed by Lyxor to prioritize investor liquidity. Thus an investor without a strong preference for liquidity is likely to prefer investment in the main fund.

For their part, hedge fund managers have several incentives to join a separate account platform: the most obvious is the fee revenue they collect on assets in their accounts, but the platform also serves to improve the name recognition and reputation of the fund. Smaller funds with weaker reputations thus have stronger incentives to join a platform than larger, more established funds. At the same time, Lyxor requires significant performance history from prospective funds and targets larger, more respected hedge funds for inclusion. Because both prospective fund managers and Lyxor know that separate accounts are subject to larger flows than traditional hedge funds, funds that successfully create a separate account on the platform are likely to be those with relatively liquid assets.

3.3.2 Separate Account and Main Fund Data

Lyxor publishes performance for all funds on the platform each week. While the weekly report contains current and historical NAVs only for funds that are live on the platform at the time of the report, we have compiled a historical dataset containing defunct fund performance as well. The survivorship and other biases associated with hedge fund returns, discussed in Fung and Hsieh (2000) and elsewhere, are not of concern in our study because we care only about the difference between matched separate account returns and associated main fund returns rather than the cross-sectional properties of the returns to the funds themselves.

Our raw dataset contains weekly historical returns for 291 Lyxor accounts between 2002 and 2010. After removing funds with less than two years of history, we have 218 Lyxor accounts. We augment these NAVs with AUM, fee, manager, and strategy information from fee documents and past weekly reports. We also add information about the benchmark main fund associated with the account. Figure 3.1 shows the total AUM over time of the 291 funds in our uncut sample. Lyxor assets under management grew rapidly until 2007, in which the assets in the accounts in our sample surpassed \$10 billion. Lyxor experienced large outflows during the financial crisis of 2007–2008, but recovered in 2009 to near 2007 levels.

A number of separate accounts on the Lyxor platform liquidated during the financial crisis and in many cases we observe dramatic poor performance during the final months before exiting the platform. Because these wind-down observations represent small sums of money and outlier returns, we exclude observations for months in which fund AUM was less than \$2.5 million at the end of each fund's reported data series. The removal of these wind-down observations, mostly present during the financial crisis, has the effect of attenuating our results.

For each fund, Lyxor provides investors with a document describing the fund's history, investment policy, and information about the associated main fund. For each separate account we find the main hedge fund identified in the investor documents in our hedge fund database, which consists of a merge of five large commercial hedge fund datasets: TASS, HFR, CISDM, Barclay Hedge, and Morningstar. These datasets were merged using fund characteristics, manager names, and return information.

Because separate account information is reported as of the close of trading each Tuesday and main fund information is reported at the end of each month, the return periods for the separate accounts may not perfectly match those of the main funds. Therefore return period ends may differ by as much as two trading days (if the month ends on a Thursday or Friday). We mark Lyxor month-end dates as the closest reported return to the calendar month-end date, including weekends. If the two Lyxor reported returns are equidistant from the month-end, we use the Lyxor date that precedes the month-end. In all cases, the Lyxor reported dates we use are within a week of the calendar month-end date. After applying all filters, we have a total of 7,171 monthly observations matched between the Lyxor and main fund datasets representing 135 fully vetted fund matches.

Table 3.1 summarizes the fund characteristics of our separate account sample and the associated main funds. Statistics for the universe of funds (not including funds of funds) reporting during the sample period are included for comparison. Both main funds and Lyxor separate accounts charge typical hedge fund management and incentive fees, approximately 20% for the incentive fee and around 1.5% for the management fee. There are no lockup periods or fundwise minimum investment amounts for the separate accounts while a minority of main funds in our sample do report initial lockup periods (most frequently of a year) and all have minimum investment amounts, often around a million US dollars. Redemption notice period is an important share restriction that determines how far in advance an investor must request a redemption before it is available. In the main funds a typical notice period is between 30 and 90 days. The notice period for Lyxor separate accounts is uniform at the significantly lower 2 trading days. When compared with the universe of hedge funds, the main funds with accompanying separate accounts in our sample have shorter lockup periods on average (0.57 months versus 3.49). Since funds self-select onto the platform, those with very illiquid assets and/or high share restrictions are unlikely to have platform separate accounts. Funds on the platform are also somewhat larger than the average hedge fund in the universe, consistent with Lyxor's goal of attracting large and reputable funds.

The frequency of redemption windows for the main funds in our sample is typically either monthly and quarterly. Lyxor separate accounts, on the other hand, have redemption windows once a week, offering investors much more frequent opportunities to redeem, invest, or rebalance. High watermark and leverage variables are not reported separately for the separate accounts because they operate the same way the main funds do in these respects. Since fund characteristics are fairly uniform across the Lyxor platform accounts (or match the associated main fund), much of our analysis will use the characteristics of the main fund, as reported by the data sources used to construct the merged main fund dataset. Lyxor accounts are typically significantly smaller than their main fund equivalent, with the mean account AUM around 66 million dollars (versus the mean of 699 million in the associated main funds).

Figure 3.2 shows the Lyxor and main fund performance of two example funds. In the first case, the Lyxor account return closely matches that of the associated main fund; in the second case there is a greater performance difference. These examples show the typical high correlation between the main fund and its associated separate account, but also illustrate the performance drag separate accounts in our sample suffer.

To examine how closely separate accounts mimic the main funds, we compute the return correlation and tracking error, fund by fund, between the main fund and separate account. Our measure of tracking error is the standard deviation of the difference in monthly return. The mean and median correlations across funds are 87.3% and 90.1%, respectively. The minimum and maximum correlations in our sample are 60.2% and 96.2%. The mean tracking error is 1.4%, with a median of 1.3%.

3.4 Results

3.4.1 Fund-Level Tests

Getmansky, Lo, and Makarov (2004) and others suggest that smoothing induced by managerial reporting discretion or stale prices will have a moving average time-series structure. Therefore to examine the smoothing of main funds and separate accounts, we use a moving average time-series model. We compute the moving average coefficients fund-by-fund, average over funds, and report our results in Table 3.2. We check moving average models of up to order three.

In Panel A of Table 3.2 we see that Lyxor funds exhibit an average MA(1) coefficient of 0.121, while the associated main funds have an average MA(1) coefficient of 0.182. The difference is statistically significant (t=4.35). Note that the moving average smoothing is economically meaningful in both cases, but the magnitude is smaller among the Lyxor accounts. Referring back to our discussion of Hypothesis 3.1, we find support for the notion that there is significant smoothing in the main hedge funds that is not present in the matching Lyxor accounts. We also note that the majority of total smoothing appears in both the Lyxor account and main hedge fund returns, so asset-level serial correlation is also a contributing factor to fund-level return characteristics.

In Panel B of Table 3.2 we see the coefficients from an MA(2) model. The first moving average coefficient for both the main fund and separate account is again significant (t=8.62 and t=5.88, respectively), but the second lags are not individually significant (t=1.28 and t=-0.54). Similarly in Panel C we see that none of the higher order coefficients in the MA(3) model is statistically significant. Using the Bayesian Information Criterion, we find that an MA(1) model is appropriate for most of the sample funds—a fact corroborated by statistical insignificance of higher order terms in panels B and C. We therefore restrict our attention to the coefficient in an MA(1) model and use this coefficient as our measure of smoothing.

If we let δ denote the moving average smoothing associated with the underlying hedge fund assets and their evaluation in the context of the Lyxor separate accounts and let γ denote the smoothing associated with the main funds but not the separate accounts (this portion of smoothing may be due to managerial manipulation of reported returns, for example) then we have $\theta_1^s = \delta$ and $\theta_1^m = \delta + \gamma$, where θ_1^s and θ_1^m are the estimated MA(1) coefficients of the separate account and main fund returns, respectively. Then we can express the proportion of hedge fund smoothing in the main funds that is due to managerial smoothing and other features unique to the main fund as

$$\frac{\gamma}{\delta + \gamma} = \frac{\theta_1^m - \theta_1^s}{\theta_1^m} \tag{1}$$

Using the estimated coefficients in Panel A of Table 3.2, we conclude that approximately 33% (=1-0.121/0.182) of hedge fund smoothing appears to be attributable to managerial reporting while 67% is due to the underlying asset and strategy characteristics.

This finding addresses a long-standing question in hedge fund research. Hedge fund returns have long been known to be serially correlated and researchers have found evidence that managers may use their discretion in reporting to distort reported returns. Nevertheless, it has been difficult to determine the extent to which this discretion is the driving force behind the serial correlation. As Getmansky, Lo, and Makarov (2004) point out, the nature of the underlying assets held by hedge funds may induce smoothed reported returns even if those returns are computed by a party with no incentive to distort returns. We find evidence for both types of smoothing, with the asset induced contribution approximately twice that of managerial discretion. One possible implication for investors or regulators is that modifying the return calculation and reporting method for hedge funds such that the manager has no discretion could eliminate a significant portion, but not all, of the observed serial correlation.

Turning our attention to the performance impact of share restrictions, Table 3.3 reports summary statistics of main fund returns and the corresponding matched Lyxor returns. Statistics are computed fund by fund and then averaged across funds. In this test, the performance criteria (average and median return) show support for Hypothesis 3.3: The Lyxor accounts underperform their associated main funds by 0.144% monthly, which corresponds to an annualized performance difference of 1.7%. This difference is statistically significant using a t-test (t-statistic=7.46) and also the nonparametric Wilcox test (p-value=0.000). We also notice that the overall volatility of the Lyxor accounts is similar to that of the main funds. Lyxor account returns are more negatively skewed and display greater kurtosis than the associated main fund returns. If the difference in performance between main funds and their separate accounts is driven by costly outflow events faced by Lyxor accounts but not by the main funds, large redemptions could create a negative tail in the Lyxor performance distribution consistent with our data.

Since the main funds and matching separate accounts have the same assets, smoothing distortion induced by the manager in the main fund should not influence the long-term performance of the fund. Ultimately, reported returns must correspond to the realized value of the assets in the portfolio. Thus, under the null hypothesis that there is no significant cost to the dramatically higher fund-level liquidity faced by Lyxor investors (and consequently more volatile flows on the Lyxor platform), the average returns for the main fund should equal that of the separate accounts. On the contrary, these results suggest that high fund-level liquidity is associated with a significant long-term performance difference. This performance difference could be the result of direct costs imposed on the fund because of forced trading of illiquid assets, for example. Alternatively, the costs of liquidity could be less direct, such as the inability of the manager to maintain the desired portfolio continuously because of the time it takes to adjust the portfolio. Either way, these results support the argument that tight hedge fund share restrictions are justified as a method for improving fund performance. While not all investors have the luxury of investing in funds with low fund-level liquidity, it appears that those who can do so benefit from it.

3.4.2 Cross-Sectional Regression Results

We are interested in identifying what fund characteristics are associated with a greater degree of managerial smoothing in the main funds. Hypothesis 3.2 suggests that increased share restrictions may give the manager greater ability to smooth reported returns, since redemptions (which require accurate and up-to-date net asset values) are less frequent. One approach is to compute the difference in smoothing, fund by fund, between the main fund and separate account and regress it on the fund characteristics. Because the Lyxor accounts have many characteristics that are homogeneous across funds (for example redemption frequency, minimum investment, and redemption notice period), we regress the MA(1) differences between the main fund and the Lyxor account on the main fund's characteristics.

We estimate the following cross-sectional regression:

$$\begin{aligned}\theta_{i,dif} =& \alpha_i + \beta_1 \text{LiquidationTime}_i + \beta_2 \text{MinInvest}_i + \beta_3 \text{RedemptionFreq}_i \\ &+ \beta_4 \text{AumRatio}_i + \epsilon_i\end{aligned} \tag{2}$$

where $\theta_{i,dif}$ is the difference between the MA(1) coefficients estimated for fund i in the main fund and in the Lyxor separate account. In our data the lockup period length and redemption notice period length are highly correlated, so we add the two together to get a measure of total time to liquidation for main fund i, LiquidationTime_i, in years. MinInvest_i is the minimum permitted investment in main fund i and RedemptionFreq_i is the length between redemption windows, in years. AumRatio_i represents the ratio of the AUM of separate account i to the average AUM of main fund i. The governance quality index of Ozik and Sadka (2015) and average AUM of the separate account and main fund were found to be statistically insignificant in this and later tests and are omitted for brevity.

Table 3.4 reports the results of these regressions. We find that the total liquidation time, a measure of fund-level illiquidity, is positively and significantly associated with greater moving average smoothing coefficients in the main fund compared to what we find in the separate account. In specification (5), an additional year of total time to liquidate implies an increase in the smoothing of the main fund of 0.126 relative to that of the separate account (t-statistic=1.91). Other variables do not significantly explain the moving average coefficient difference between main funds and their separate accounts after liquidation time is taken into account.

Recall that we have mentioned two explanations for smoothed hedge fund returns: (1) The underlying assets of the fund have serially correlated returns—at least, as computed using the method managers and Lyxor use to value them—and (2) managers use valuation prerogative to induce serial correlation in reported returns to improve the fund's appearance or to maximize their fees. High fund-level liquidity makes any type of manipulation of reported returns more difficult since subscriptions and redemptions must be performed at a NAV that is reported to the investors, which makes reporting a different NAV to data vendors and investors a risky proposition. Further, manipulating the NAV used for redemptions and subscriptions would lead to wealth transfers between investors subscribing/redeeming shares and those remaining in the fund, a situation the manager likely wants to avoid. While manipulation is possible in the main fund in any month, it is riskier or more difficult to smooth returns if there are flows. Our observation in Table 3.4 that funds with higher share restrictions have a greater share of "managerial" smoothing supports Hypothesis 3.2.

One may ask whether the assets in the main fund and the Lyxor account differ to a greater degree for funds with lower fund-liquidity. After all, we expect barriers to liquidity to be associated with lower asset liquidity if the premise of Hypothesis 3.4 is satisfied. By agreement, though, the Lyxor account trades in the same assets and at the same time and same price as the main fund. The differences between the two should be driven by the differences in flows and the consequent impossibility of making perfectly parallel trades. For example, if positive flows in the Lyxor account are held as cash by the manager until a (relatively infrequent) positive flow is made in the main fund and if the returns of the fund's underlying assets are autocorrelated, then the extra cash in Lyxor accounts would decrease those accounts' serial correlation relative to the main fund. Of course, this effect is limited to time periods in which the Lyxor accounts experience different flows from the main fund and the Lyxor manager consequently holds assets as cash. It is also limited by the proportion of assets invested in cash-like securities.

We do not necessarily expect a change in serial correlation associated with outflows on Lyxor (though there may be a difference in overall performance). In reference to the overall discrepancy in smoothing in Table 3.2 of 0.182 versus 0.121, Lyxor accounts would need to have a very large proportion of total assets ($1 - 0.121/0.182 = 0.335$), on average, in cash to fully explain the difference. We would further need to multiply this ratio by the possible proportion of the time Lyxor accounts have had cash inflows but in which the managers are unwilling to invest the cash to estimate the true cash fraction. Overall it seems unlikely that this effect accounts for the difference in smoothing between Lyxor and the main fund.

To examine the effect of increased fund liquidity on performance, for each fund we compute the difference in average returns between the main fund and the associated separate account:

$$\tilde{r}_{i,dif} = \text{Main Fund } \bar{r}_i - \text{Lyxor } \bar{r}_i \tag{3}$$

where Main Fund $\bar{r}_i$ denotes the time-series average return to the main fund i and Lyxor $\bar{r}_i$ denotes the corresponding time-series average return to that fund's associated Lyxor separate account. Table 3.5 reports the results from the following cross-sectional regression:

$$\begin{aligned}\tilde{r}_{i,dif} =& \alpha_i + \beta_1 \text{LiquidationTime}_i + \beta_2 \text{MinInvest}_i + \beta_3 \text{RedemptionFreq}_i \\ &+ \beta_4 \text{AumRatio}_i + \epsilon_i\end{aligned} \tag{4}$$

Each of these fund characteristic variables pertains to the main fund. As in Equation (2), we combine the lockup period length and redemption notice period length to get the total liquidation time for each fund. Also included are the minimum investment, redemption frequency, and the ratio of the average AUM of the separate account and main fund. It should be remembered that investors in the separate accounts must pay the same management and incentive fees as the main-fund investors, so these variables are not relevant to a cross-sectional difference regression.

Table 3.5 reports results of these regressions. Total liquidation time is a significant positive predictor of the performance difference between a fund and its associated separate account. In specification (5), an additional year of time to liquidation implies a monthly performance difference between the main fund and separate account of 0.186% (t-statistic=2.08). Both components of total liquidation time (lockup period length and redemption notice period length) are forms of share restrictions. We may interpret a relation between these variables and the performance difference between the main fund and separate account as higher share restrictions protecting investors in the main fund from the performance reducing effects of fund-level liquidity. Another interpretation is that funds impose higher share restrictions in cases where they are more susceptible to the costs of fund liquidity. This argument supports Hypothesis 3.4.

Additionally, in specification (5), the ratio of the average size of a Lyxor separate account to the average size of the associated main fund negatively and significantly predicts the size of the performance gap. A unit change in this ratio causes a -0.053 change in average monthly performance (t-statistic=-2.13). Recall that one of our primary explanations for the performance gap between main funds and separate accounts is that the reduced barriers to investor flows in the separate accounts impose greater trading costs on the separate account. If there is a component of these trading costs that is fixed or is proportionally reduced with scale, then this result is expected: when there is a larger gap in size between a main fund and its separate account, the costs associated with investor liquidity are proportionally greater for the separate account than when the size difference is small.

To summarize our fund characteristics analysis, we find that funds with less fund-level liquidity, as measured by the total liquidation time faced by a new investor, have greater serial correlation difference between a main fund and its associated separate account. Artificial smoothing, introduced by choice of valuation and reporting methodology, can easily be applied to main funds that do not expect flows due to high share restrictions. Separate accounts, however, must always expect flows, so Lyxor is less likely to use those valuation techniques in its separate accounts. Thus we find a greater serial correlation gap for funds with high share restrictions. On the performance side, hedge funds that protect investors in the main fund from the costs of flows by restricting liquidity through share restrictions have a greater performance advantage over matching separate accounts than those that do not.

3.4.3 The Cross-Section of Flow Volatility

The fund characteristics used in Section 3.4.2 to explain the cross section of differences between main fund and separate account returns can be thought of as measures of the manager's *ex-ante* view of the costs of fund-level liquidity. Managers who believe that low barriers to investor subscriptions and redemptions would be costly may choose more onerous share restrictions. Similarly, a manager desiring to engage in smoothing of reported returns might choose high share restrictions. Equivalently, managers with greater share restrictions may find smoothing of reported returns to be easier. In both cases we consider managerial beliefs about the fund prior to its inception. These choices may be driven by beliefs about the costs of subscriptions and redemptions, or beliefs about the magnitude of these events. We may then ask whether large flows do affect performance and whether they affect the smoothing of hedge fund returns.

Perhaps the most natural measure of the magnitude of flows is the standard deviation of flows, where dollar flows each period have been normalized by dividing by fund AUM at the beginning of the period. In our sample, this standard deviation has some outliers due to a few fund-months with low AUM. To mitigate the effect of these outliers we modify our normalization by dividing dollar flows not by the prior period AUM but by the time-series average AUM for that fund. Specifically,

$$\mathrm{NormFlow}_{i,t} = \frac{\mathrm{AUM}_{i,t} - \mathrm{AUM}_{i,t-1}(1 + r_{i,t})}{\frac{1}{T}\sum_{s=0}^{T} \mathrm{AUM}_{i,s}}. \tag{5}$$

This measure looks both forward and backward in the time series in order to compute the denominator. For this reason we do not use it in our time-series analysis. However, this measure does not suffer from the outlier problem, so we use it in our cross-sectional analysis.

Panel A of Table 3.6 reports results of the cross-sectional regression of the average performance difference between the each main fund and its separate account on the standard deviation of normalized flows of the separate account and the main fund. Since most of the main funds have relatively stable flows, we first examine the effect of the standard deviation of flows in the Lyxor account on the performance difference in specification (1). Higher flow standard deviation on Lyxor is positively associated with a larger performance difference between the separate account and main fund (coef=0.447%, t=2.22). Adding the standard deviation of flows in the main funds in specification (3) we see that a greater standard deviation of the normalized flows in the Lyxor separate accounts predicts a greater main fund/Lyxor performance differential (0.505% monthly, t=2.33). Additionally we find that if the main fund also has a greater standard deviation of flows, the performance differential is smaller by -0.539% (t=-1.95) and the R^2 of the specification including both variables is noticeably higher (7.1%). Taken together, these results suggest that if the disparity between flows in the separate account and main fund is greater, the performance difference between the main fund and separate account will also be greater. This could be the case if, for example, flows impose direct costs on the fund that drag down performance. Alternatively, high flows could affect performance indirectly by preventing the manager from achieving the goal of exactly matching the portfolio composition of the main fund.

In Panel B of Table 3.6 we report the results of a similar regression in which we use the standard deviation of flows to explain differences in moving average smoothing process coefficient between the main fund and its separate account. If it is the case that the smoothing observed in the main fund reported returns is facilitated by low incidence of flows (within the limits of the share restrictions), then we might expect flows to predict the smoothing differential. Empirically we find no evidence that the volatility of fund flows explains smoothing in the cross-section. In short, we find that the most reasonable explanation for the difference in MA(1) coefficient between main funds and separate accounts is manager discretion in how to calculate and report returns.

3.4.4 Time Series Return–Flow Relations

Now we turn to the examination of the relation between fund returns and subsequent fund flows separately for the Lyxor separate accounts and for the main funds. Fund flows often show significant autocorrelation, so we control for lagged fund flows where appropriate. Fund flows are calculated at the monthly frequency, using

$$\text{Flow}_t = \frac{\text{AUM}_t - \text{AUM}_{t-1}(1 + r_t)}{\text{AUM}_{t-1}}, \tag{6}$$

consistent with Sirri and Tufano (1998) and others.

We regress fund flows to the separate accounts on past flows and past performance using a fund fixed-effects regression.

$$\text{Flow}_{i,t} = \sum_{j=1}^{4} (\beta_j r_{i,t-j} + \gamma_j \text{Flow}_{i,t-j}) + \epsilon_{i,t} \quad (7)$$

where $\text{Flow}_{i,t}$ is the flow into Lyxor account i during month t (i.e., from time $t-1$ to time t) and $r_{i,t}$ is the return to account i during the month from $t-1$ to t.

Table 3.7 reports results using Lyxor separate account data. A positive coefficient for a return variable indicates positive flow in response to past fund performance, that is, investors chasing performance at the monthly frequency. We find a positive relation between past returns and future flows for up to three monthly lags (coefficients are 1.640, 0.428, and 0.307). We cluster standard errors by fund and correct for heteroskedasticity as in Beck and Katz (1995). The resulting t-statistics 14.40, 3.68, and 2.77, respectively, in specification (4) and the R^2 is 6.5%. These results indicate that investors chase performance on the Lyxor platform in the very next month and that those flows take several months to completely filter into the account. Moreover, given the positive serial correlation observed in Table 3.2, we conclude that investors who chased performance at the monthly frequency during our sample benefited, on average, from doing so. At the same time, the performance chasing behavior did not cause enough flow-driven trading to affect prices and drive this serial correlation to zero. We also find strong positive autocorrelation in flows up to four months back.

Next we examine the return–flow relation in the main funds by estimating the model in Equation (7) using data from the main funds. Table 3.8 reports these results and shows a different picture. Flow autocorrelation is not significant in the first lag and is negative in other lags. Once lagged returns are not significantly related to flows, which would be consistent with investors being either unable to observe returns or invest in the subsequent month. The two-month lagged return is significant in specifications (2) through (4), but the R^2 for each specification is near zero. In short investors appear to be unable to chase performance at short horizons and to the same degree among main funds that they do in the separate accounts. These results support Hypothesis 3.5, that the greater ease of moving money from one fund to another on the Lyxor platform leads to stronger performance-chasing behavior on the Lyxor platform than among the associated main funds.

3.4.5 Main Fund Return Prediction

Since separate account returns are calculated by a third party and reported almost immediately while main fund returns are calculated by the fund itself and reported relatively infrequently, there may be investors who have access to updated Lyxor performance information but do not yet know how the main funds have fared. One might then ask whether it is possible to use available separate account returns to predict the subsequent performance of the main fund before the contemporaneous main fund returns become available. For any investors seeking to time flows in the main fund but who lack up-to-date performance and flow information on the main fund, this information could be useful. In Table 3.9 we examine the degree to which lagged Lyxor returns explain main fund returns (in isolation of past main fund information). The fund fixed-effect regression equation is

$$\text{Main Fund } r_{i,t} = \sum_{j=1}^{4} \left(\beta_j \text{Lyxor } r_{i,t-j} + \gamma_j \text{Lyxor Flow}_{i,t-j}\right) + \epsilon_{i,t} \tag{8}$$

where Main Fund $r_{i,t}$ is the return in the main fund i in month t and Lyxor $r_{i,t}$ is the associated separate account return in month t. Additionally, Lyxor $\text{Flow}_{i,t}$ is the flow to the separate account in month t. We see that a single lag of Lyxor returns does indeed explain main fund returns—in specification (1), the first lagged Lyxor return has a coefficient of 0.145 (panel-corrected t-statistic=9.53). The ability of separate account returns to predict future main fund reported returns means that an investor with access to Lyxor returns (which are published almost immediately) would have information about the returns that ultimately will be published in the next period for the main fund. If the main fund's share restrictions permit it, such an investor could better time the movement of cash into or out of the main fund.

3.4.5.1 Time Series of the Performance Gap

If the main fund returns are essentially a smoothed version of the separate account returns, we would expect that shocks to the performance gap between the main fund and separate account would subsequently be reversed. It is also possible that the performance gap may have power to predict future main fund returns. We examine these questions by studying the time-series of returns for each fund/separate-account pair.

We first test whether the performance difference between the main fund and separate account predict subsequent returns in the main fund. The regression equation for a single fund is

$$r_t^m = \alpha + \beta_1(r_{t-1}^m - r_{t-1}^s) + \beta_2 r_{t-1}^m + \epsilon \tag{9}$$

where r_t^m is the return to the main account and r_t^s is the return to the separate account in month t. If the performance gap predicts main fund returns, we would expect β_1 to be statistically different from zero. We run this regression once for each of our 135 fund/account pairs and find that for funds (5.9%), β_2 is negative and significant and for 3 funds (2.2%) it is positive and significant. Overall we do not find evidence that the gap predicts overall main fund returns.

We also test whether the performance gap is directly reversed in the following period, consistent with main fund returns being a smoothed version of separate account returns. The regression equation for this test (for a single fund) is

$$(r_t^m - r_t^s) = \alpha + \beta_1(r_{t-1}^m - r_{t-1}^s) + \beta_2 r_{t-1}^m + \epsilon. \tag{10}$$

We would expect β_1 to be negative and significant if the performance gap is partially reversed in one period. We run this regression separately for each fund pair in our sample. For 69 of our 135 funds (51%), β_1 is negative and statistically significant at the 5% level while it is positive and significant for only 5 funds. Overall it is clear that the performance gap between the main fund and separate account is at least partially reversed within a month in many cases, though the effect is small compared with the overall variation in fund returns over time.

3.4.6 Mechanisms of Discrepancies

Lyxor separate accounts are managed concurrently with the associated hedge funds by the same managers and have the same objectives. We may then ask why we observe return differences at all, either in the short or long run. Here we discuss some possibilities.

3.4.6.1 Flow-Timing Costs

By agreement with Lyxor, separate account and main fund assets trade together to prevent front-running. However, this does not imply that the trades will always mirror each other perfectly nor that the assets in separate accounts will always match those of the main funds exactly. Investors in the separate accounts are free to make redemptions and subscriptions on a weekly basis while investors in the main fund have no such flexibility. Potentially costly trading to meet liquidity demands will therefore be much heavier in the separate account. In the presence of any asset-level illiquidity and large flows, the manager may not be able to trade to a matching portfolio in the separate account immediately. In the case of large separate account inflows, for example, the manager would need to leave a proportion of new money either in cash or in more liquid portfolio assets until the matching assets can be purchased. In the case of large outflows, the manager may need to sell liquid portfolio securities while temporarily being unable to sell those that are less liquid. These cases are in the minority and involve a small proportion of the funds' assets[4].

Cash management is another possible source of performance difference. Hedge fund managers may have a number of strategies for holding cash while waiting to trade it in their fund or while preparing for redemptions. These options may include the temporary purchase of liquid but risky securities. On the other hand, Lyxor manages all the cash for the separate accounts by investing in US treasury bills. The proportion of assets held in cash may differ between main funds and separate accounts as well. After large separate account subscriptions or perhaps in anticipation of large redemptions, the manager may hold more cash in the separate account than she needs to in the main fund.

As researchers we do not observe whether long-term performance difference is attributable to the manager being forced to hold more assets in cash, being unable to take positions in some illiquid securities, or having to bear greater transactions costs in the separate account. These can all be viewed, however, as costs imposed on separate account investors in exchange for greater liquidity. If a manager is unable to perfectly match the target portfolio in the separate account because of liquidity concerns, we can say that the difference between the realized and desired returns are the costs of fund-level liquidity.

3.4.6.2 Reporting

Reporting differences between the main fund and separate account may induce short-term, but not long-term, return differences. We previously mentioned that a manager's method for valuing assets in the main fund may differ from those used by Lyxor in the separate account. If the manager's method uses relatively stale prices, for example, we would observe greater positive serial

[4]Thanks to Lionel Paquin and Stefan Keller from Lyxor Asset Management, who provided this insight about Lyxor's separate accounts.

correlation in the main fund than in the separate account without an associated difference in the long-term mean return.

3.4.6.3 Holdings Differences

We may ask whether manager inability to match portfolio holdings in the Lyxor account induces the manager to purchase substitute, non-mimicking assets in the separate account. If so, this could drive the differences in performance and serial correlation that we observe. If a manager systematically held more illiquid assets in the main fund than in the separate account, for example, the underlying assets of the latter may induce less serial correlation in separate account returns. If the illiquid assets in the main fund further earn an illiquidity premium, we might expect a performance difference between the two as well. Importantly, in this case the two effects—poor performance in separate accounts and less serial correlation in separate accounts—would be linked. Because the manager's inability to hold illiquid assets in the Lyxor account is a result of the difference in flows in the two accounts, we would also expect this effect to be strongest for funds with relatively low flows in the main fund and high flows in the separate account.

Referring back to Panel A of Table 3.6, we we do indeed see a greater performance gap as funds have large flows in the separate account and relatively low flows in the main fund. However, looking at Panel B, we do not observe a relation between serial correlation differences and differences in flow magnitude, as the long-term portfolio differences explanation would suggest. The cause of the serial correlation gap is not directly tied to flows.

Moreover, if systematic holdings differences between separate accounts and main funds were important drivers of our results, we would expect main funds to load more heavily on measures of illiquidity than separate accounts do. We do not find evidence that this is the case. We regress main fund and separate account returns on the Pastor-Stambaugh (2003) liquidity factor with controls for the market excess return, SMB size factor, HML value/growth factor, and MOM momentum factor. The regression equation is

$$r_{i,t} = \alpha_i + \gamma_1 \mathrm{LIQ}_t + \gamma_2 \mathrm{MktRf}_t + \gamma_3 \mathrm{SMB}_t + \gamma_4 \mathrm{HML}_t + \gamma_5 \mathrm{MOM}_t + \epsilon_{i,t} \tag{11}$$

We perform this regression fund-by-fund and then compute the average value of γ_1 across funds to measure the liquidity risk exposure. Using only the first three factors as controls, we find that the average γ_1 for main funds is 0.042 and for separate accounts it is 0.049. The difference is insignificant, with a t-statistic of -0.857, and the sign is the opposite of the prediction would be if managers avoided illiquid assets in separate accounts. Adding the momentum factor, the main fund average γ_1 becomes 0.030, versus a separate account coefficient of 0.033. Again, the difference is insignificant with a t-statistic of -0.454.

We conclude that portfolio holdings differences between separate accounts and main funds are small and/or short-lived in nature. This is consistent with the objective of the separate accounts to track the associated main funds by mimicking their holdings. While flows are costly to separate account investors, they do not appear to force managers to maintain different assets in the main fund and separate account, nor do they directly contribute to the serial correlation gap. This evidence suggests that the difference in serial correlation between main funds and separate accounts is a function of portfolio valuation and reporting practices.

3.4.6.4 Flow-Induced Autocorrelation

Large flows to a fund naturally lead to trading within the fund as the manager takes positions to utilize inflows or liquidates positions to satisfy outflows. Since fund flows are persistent and predictable, repeated flow-driven transactions can induce serial correlation in a fund's returns (as pointed out by Lou, 2012). Since the separate accounts in our study permit much larger proportional flows than the main funds, we may ask whether the flow-induced trading in the separate accounts affects the serial correlation of assets in the separate account and main funds. If so, we would expect to see an increase in serial correlation in the main funds during the period when the separate account is active relative to the period before the creation of the separate account. Of our 135 fund pairs, 99 main funds have 15 or more observations reported for the main fund before the creation of the separate account. For these 99 funds we have an average of 62 monthly observations per fund before the beginning of their separate accounts and 55 observations during the separate account period. We compute the MA(1) coefficient for each fund during the pre-separate account period and again during the separate account period. The mean MA(1) coefficient before the separate accounts was 0.092, compared to 0.203 during the separate account period (t-statistic=3.44).

The increase in serial correlation in the main funds over time is consistent with the explanation that flow-based trading in the separate accounts drives up serial correlation of these funds' assets. However, time-period and fund life-cycle effects could easily cause this same result (if serial correlation across all hedge funds was higher in the second sample than the first or if older, more established funds have greater serial correlation than those that are new). To test whether it is the separate accounts or other factors that caused the increase in serial correlation, we construct a matched hedge fund sample and examine the changes in serial correlation over time.

For each main fund in our sample, we identify a similar "matched" fund that is not associated with a Lyxor separate account. The matched fund reports to the same hedge fund database as the sample fund, has the same strategy classification, is live during the period when the separate accounts were live, and came from a different fund family than the sample fund. For a given sample fund, we select the best match from the permissible candidates by identifying the match fund that had its inception date closest to the sample fund's inception. In cases where there are ties, we select the match fund with the AUM that was most similar three months after inception to the sample fund's AUM three months after its inception. We do not use the first available AUM because initial funding flow timing often distorts the first few AUM observations.

If the creation of the separate accounts caused the increase in serial correlation observed in our sample's main funds, then the same increase should not be present in the funds in our matched sample (which had no separate accounts). For each fund in the matched sample, we compute the MA(1) serial correlation before and during the life of the separate account associated with the matching main fund from our sample. We find that the matched sample MA(1) coefficients rise, on average, from 0.115 to 0.165 (t=4.70) even though these funds are not associated with a separate account. We then construct the difference between the main fund and matched fund serial correlation in both time periods

$$\theta_{i,0}^{dif} = \theta_{i,0}^{m} - \theta_{i,0}^{*} \tag{12}$$

$$\theta_{i,1}^{dif} = \theta_{i,1}^{m} - \theta_{i,1}^{*} \tag{13}$$

where $\theta_{i,0}^{m}$ represents the MA(1) coefficient for the main fund i from our sample before the period where its Lyxor account was started. $\theta_{i,0}^{*}$ is the MA(1) coefficient for the matched sample fund that never had a separate account, over the same period. $\theta_{i,1}^{m}$ represents the MA(1) coefficient for main fund i from our sample during the period when its Lyxor account was live. If creation of separate accounts increases the serial correlation in main funds, then $\theta_{i,0}^{dif}$ should be greater than $\theta_{i,1}^{dif}$. The

t-statistic for equality of these two coefficients is 0.89 (p=0.375). Since the serial correlation in funds that never had a separate account rises at the same time and by a statistically equal amount as the funds in our sample, we conclude that life-cycle or time-period effects are the primary drivers in the increase in serial correlation in our main funds. We do not believe flow-induced trading in the separate accounts drive serial correlation in the main funds—remember that although the percentage flows in the separate accounts are large, the accounts themselves are relatively small, so flow-induced serial correlation effects are likely to be small as well.

3.4.6.5 Payoff Complementarities

Because the separate account and main fund attempt to hold the same assets, the possibility of payoff complementarities based on predictable flows exists[5]. For example, if the separate account return is low in period (t-1), investors in the main fund may anticipate outflows from the separate account and consequent flow-driven liquidations. These liquidations could depress the price of the assets of interest. Therefore main fund investors may add cash to their fund in order to cause flow-driven purchases at an artificially depressed price. This would result in superior performance in the main fund and, as an implication of Chen, Goldstein, and Jiang (2010), this relation would be especially strong for funds with illiquid assets.

If investors time flows to take advantage of these complementarities, we would expect a negative relation between main fund flows and lagged flows to lagged separate accounts. Specifically, for a given fund, we perform the regression

$$\text{Flow}_t^m = \alpha + \beta_1 \text{Flow}_{t-1}^s + \beta_2 \text{Flow}_{t-1}^m + \epsilon_t \tag{14}$$

where Flow_t^s is the flow to the separate account in period t and Flow_t^m is the flow to the associated main fund in period t. If strong payoff complementarities exist for that fund, we would expect a negative an significant β_1 coefficient. From our sample of 135 funds, this is only the case for 2 funds (at the 5% level). On the other hand, there are 30 funds for which β_1 is positive at the same significance level. This suggests that investors in the main fund and separate account tend rather to follow a similar, not complementary, subscription and redemption pattern. The average contemporaneous correlation between the main fund flows and separate account flows is

[5]Thanks to an anonymous referee for pointing this out.

also positive (17%). If strong flow complementarities existed, we would expect to find a negative correlation.

3.5 Conclusion

This paper answers some important outstanding questions in the hedge fund literature by taking advantage of a unique institutional setting: one in which managers run a hedge fund and a matched separate account *pari passu* so that differences between the two investment vehicles have to do with institutional structure, not underlying investment. Investors on the platform have access to immediate performance information calculated by a third party and have dramatically greater freedom to move cash with little notice. These features are used to determine (1) how much of hedge fund return smoothing is inherited from the underlying assets as opposed to managerial reporting prerogative and (2) how great is the performance reduction resulting from eliminating hedge fund share restrictions.

We compute the MA(1) coefficient in the main funds and separate accounts as our measure of smoothing. The main funds' smoothing can then be decomposed into the portion inherited from the funds' assets (67%) and the part attributable to main fund-specific considerations (33%). We may interpret the latter as the portion due to managerial manipulation/discretion. We conclude that the majority of reported smoothing is present in the underlying asset returns of hedge funds, but that there is also significant smoothing induced by the managers' reporting choices; for example, managers may modify main-fund returns so that they appear smoother or so that fees are maximized. Funds that have stricter share restrictions, such as a longer time to liquidation, appear to have greater disparity in moving average smoothing process coefficient between the main fund and separate account but fund flows themselves appear to have no effect on smoothing. One interpretation of this result is that share restrictions untie managers' hands and allow them greater freedom to report smoothed returns.

Comparing the performance of hedge funds with that of their associated Lyxor separate accounts, we find that Lyxor accounts suffer a statistically significant annualized penalty of 1.7%. This suggests that hedge funds that are not protected from frequent investor flows suffer from significant performance degradation. The magnitude of this effect appears to be related to the share restrictions of the fund such as the total time to liquidation. In addition, we find that the performance gap between the main fund and separate account is greater for funds with a greater difference in flow volatility between the main fund and separate account. Taken together, we find evidence that share restrictions do contribute significantly to hedge fund performance by inhibiting flows and that funds for which this consideration is most important (those that would suffer the most from unrestricted flows) rationally impose higher share restrictions, at least in the form of redemption notice periods or lockup periods.

Because funds with liquid underlying assets and not given to return manipulation self-select onto the separate account platform, we believe that our result in both cases forms a lower bound for that of the rest of the hedge fund universe. That is, we expect that a similar non-platform fund would experience at least as much reduction in smoothing and performance degradation as the funds in our sample if share restrictions were eliminated and reporting was done by a third party.

We also find evidence that investors on the separate account platform chase past performance from the first to third lagged month but observe no such strong effect in the main funds' flows. Hedge fund share restrictions and delayed reporting appear to be relatively effective in preventing investors from engaging in this behavior, especially at short time horizons. In fact, since Lyxor returns are available almost immediately while the same is not true for main funds, an investor with access to Lyxor returns may be able to time entrance or exit into the main funds.

In summary, the Lyxor separate account platform provides a unique laboratory in which we can measure the effect of third-party reporting, decompose smoothing into asset-induced and managerial components, and observe the costs of reducing share restrictions among hedge funds. Collectively, this paper uses a unique approach to address several important open questions that cannot be answered using conventional hedge fund data.

References

Agarwal, Vikas, Naveen Daniel, and Narayan Naik, 2004, Flows, performance, and managerial incentives in hedge funds, in *EFA 2003 Annual Conference*.

Agarwal, Vikas, Naveen Daniel, and Narayan Naik, 2011, Do hedge funds manage their reported returns?, *Review of Financial Studies* 24, 2221–2256.

Agarwal, Vikas, and Narayan Naik, 2000, Multi-period performance persistence analysis of hedge funds, *Journal of Financial and Quantitative Analysis* 35, 327–342.

Ang, Andrew, and Nicolas Bollen, 2010, Locked up by a lockup: Valuing liquidity as a real option, *Financial Management* 39, 1069–1096.

Aragon, George, 2007, Share restrictions and asset pricing: Evidence from the hedge fund industry, *Journal of Financial Economics* 83, 33–58.

Aragon, George, Bing Liang, and Hyuna Park, 2013, Onshore and offshore hedge funds: Are they twins?, *Management Science* (forthcoming).

Asness, Cliffort, Robert Krail, and John Liew, 2001, Do hedge funds hedge?, *Journal of Portfolio Management* 28, 6–19.

Beck, Nathaniel, and Jonathan Katz, 1995, What to do (and not to do) with time-series cross-section data, *American Political Science Review* 89, 634–647.

Bollen, Nicolas, and Veronika Pool, 2008, Conditional return smoothing in the hedge fund industry., *Journal of Financial and Quantitative Analysis* 43, 267–298.

Bollen, Nicolas, and Veronika Pool, 2009, Do hedge fund managers misreport returns? evidence from the pooled distribution, *Journal of Finance* 64, 2257–2288.

Chen, Qi, Itay Goldstein, and Wei Jiang, 2010, Payoff complementarities and financial fagility: Evidence from mutual fund outflows, *Journal of Financial Economics* 97, 239–262.

Fung, William, and David Hsieh, 2000, Performance characteristics of hedge funds and commodity funds: Natural vs. spurious biases, *Journal of Financial and Quantitative Analysis* 35, 291–307.

Getmansky, Mila, Andrew Lo, and Igor Makarov, 2004, An econometric model of serial correlation and illiquidity in hedge fund returns, *Journal of Financial Economics* 74, 529–609.

Hendricks, Darryll, Jayendu Patel, and Richard Zeckhauser, 1993, Hot hands in mutual funds: Short-run persistence of relative performance, 1974–1988, *Journal of Finance* 48, 93–130.

Jagannathan, Ravi, Alexey Malakhov, and Dmitry Novikov, 2010, Do hot hands exist among hedge fund managers? An empirical evaluation, *Journal of Finance* 65, 217–255.

Joenvaara, Juha, and Robert Kosowski, 2013, Geography, liquidity and fund performance: New evidence from UCITS hedge funds, Working Paper, Imperial College Business School.

Liang, Bing, 1999, On the performance of hedge funds, *Financial Analysts Journal* 55, 72–85.

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

Loudon, Geoff, John Okunev, and Derek White, 2006, Hedge fund risk factors and the value at risk of fixed income trading strategies., *Journal of Fixed Income* 16, 46–61.

Ozik, Gideon, and Ronnie Sadka, 2015, Skin in the game versus skimming the game: Governance, share restrictions, and insider flows, *Journal of Financial and Quantitative Analysis* forthcoming.

Sirri, Erik, and Peter Tufano, 1998, Costly search and mutual fund flows, *Journal of Finance* 53, 1589–1622.

Zheng, Lu, 1999, Is money smart? A study of mutual fund investors' fund selection ability, *Journal of Finance* 54, 901–933.

Figure 3.1: Lyxor assets under management

Plot of assets under management for all funds in the Lyxor separate account sample (before merging with the main hedge fund databases). Assets under management are calculated at the end of each calendar year and are reported in US dollars.

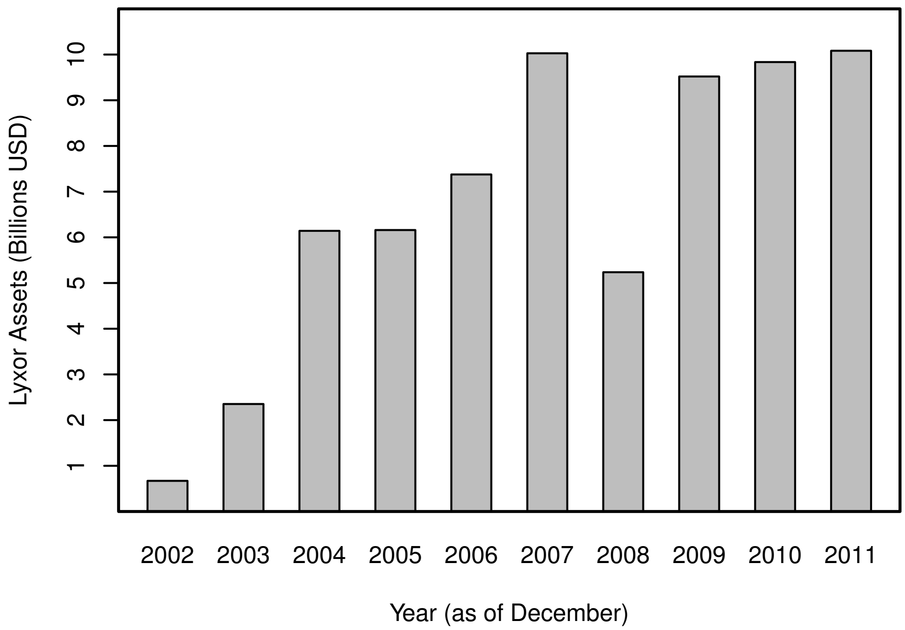

Figure 3.2: Net asset value of two representative funds: main funds versus separate accounts

Plots of the NAV of two representative funds in our sample. Each plot shows the performance of the main fund (solid line) and the associated Lyxor account (dashed line). The first example is a fund with comparatively little performance difference over time while the second has more. Initial NAVs are set to 100 so the comparison is straightforward.

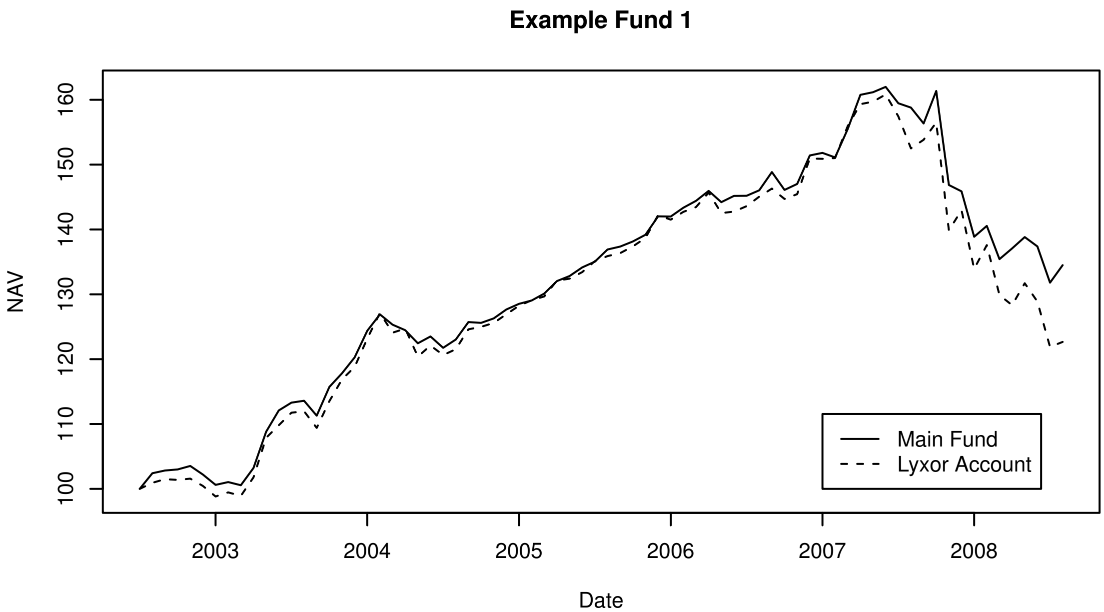

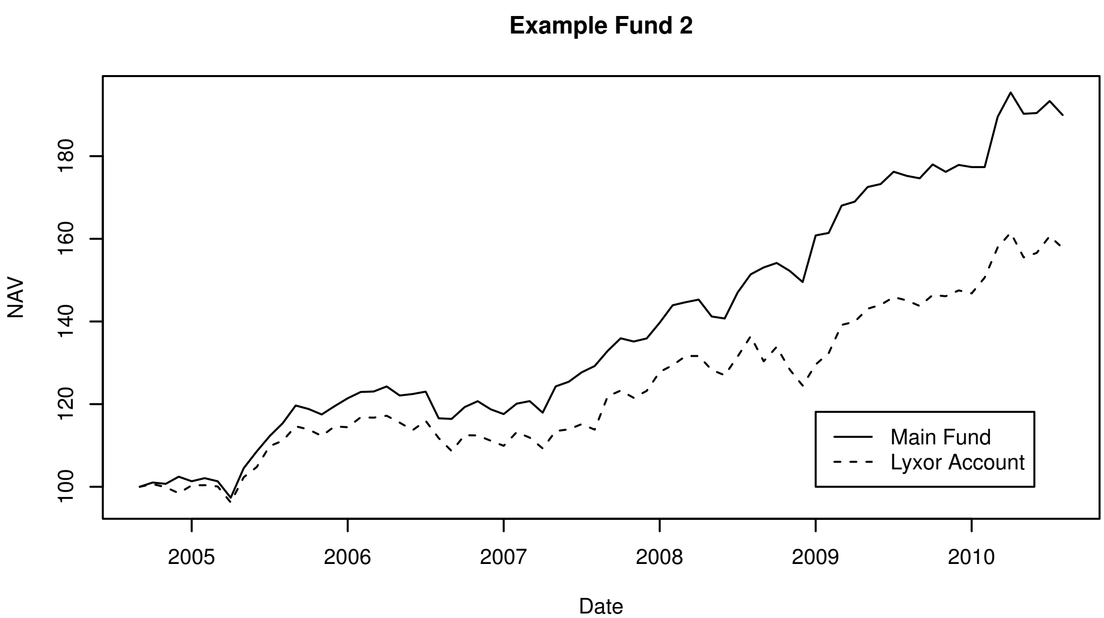

Table 3.1: Summary statistics of hedge fund characteristics

Summary statistics of the fund-level characteristics of the main funds and Lyxor separate accounts. Each Lyxor/main-fund pair is an observation. Incentive and management fees are expressed in percentages and the redemption frequency (the period of time between redemption windows) and redemption notice period are in months. Lockup period is reported in months. Minimum allowed investment is in millions of dollars. High watermark and leverage allowed are indicator variables denoting whether the fund reports using a high watermark provision and whether it reports using leverage. Separate accounts mimic the main funds in their use of leverage and high watermark provision. The variable average AUM is the average reported monthly AUM, expressed in millions of US dollars. There are 135 matched funds in this sample and the sample period is 2002-2010. Statistics for all hedge funds in the merged hedge fund universe during this time period are included for comparison.

	Main Fund		Lyxor Account		HF Universe	
	Mean	Median	Mean	Median	Mean	Median
Incentive Fee (%)	19.20	20.00	20.35	20.00	18.28	20.00
Management Fee (%)	1.52	1.50	1.58	1.50	1.45	1.50
Lockup Period (mo)	0.57	0.00	0.00	0.00	3.49	0.00
Redemption Frequency (mo)	2.00	1.00	0.23	0.23	2.36	1.00
Redemption Notice Period (mo)	1.00	1.00	0.07	0.07	1.08	1.00
Minimum Allowed Investment (MM$)	1.71	1.00	0.00	0.00	0.89	0.50
High Watermark (Yes=1)	0.85	–	–	–	0.79	–
Leverage Allowed (Yes=1)	0.61	–	–	–	0.52	–
Average AUM (MM$)	699.18	138.59	65.67	48.95	427.61	86.48

Table 3.2: Lyxor and main fund smoothing

The moving average process coefficients, fund by fund, for each fund and each separate account are estimated for models up to order three:

$$r_{i,t} = \alpha_i + \epsilon_{i,t} + \theta_{i,1}\epsilon_{i,t-1} + \theta_{i,2}\epsilon_{i,t-2} + \theta_{i,3}\epsilon_{i,t-3}$$

where $\theta_{i,j}$ represents the j^{th} MA coefficient from the time-series regression of fund or account i's returns. The mean of these fund coefficients are reported, along with their t statistics (in parentheses). Also included are nonparametric Wilcox tests for differences between main fund and separate account coefficient ranks.

Panel A: MA(1) Model

	Main Fund	Lyxor	Difference
Average θ_1	0.182	0.121	0.061
t-statistic	(9.22)	(6.28)	(4.35)
Wilcox p-value	–	–	(0.00)

Panel B: MA(2) Model

	Main Fund	Lyxor	Difference
Average θ_1	0.195	0.122	0.073
t-statistic	(8.62)	(5.88)	(4.16)
Wilcox p-value	–	–	(0.00)
Average θ_2	0.025	-0.011	0.037
t-statistic	(1.28)	(-0.54)	(2.04)
Wilcox p-value	–	–	(0.01)

Panel C: MA(3) Model

	Main Fund	Lyxor	Difference
Average θ_1	0.197	0.112	0.085
t-statistic	(8.77)	(5.12)	(5.01)
Wilcox p-value	–	–	(0.00)
Average θ_2	0.021	-0.005	0.026
t-statistic	(0.99)	(-0.23)	(1.31)
Wilcox p-value	–	–	(0.01)
Average θ_3	-0.018	-0.019	0.001
t-statistic	(-0.84)	(-0.81)	(0.07)
Wilcox p-value	–	–	(0.86)

Table 3.3: Distributional properties of Lyxor and main fund returns and flows

Summary statistics of the main hedge fund returns and normalized flows and the corresponding matched Lyxor separate account returns. Each statistic is computed fund by fund and then the cross-sectional average of the statistics is reported. Returns are in monthly percentage terms. Mean and median flows are the mean and median of the absolute value of flows, while the standard deviation is computed on signed flows. The t- and Wilcox-tests are tests performed on the cross-section of each fund statistic.

Panel A: Fund Returns

	Main Fund	Lyxor	Difference	t-statistic	Wilcox p-value
Mean	0.487	0.343	0.144	(7.46)	(0.00)
Median	0.587	0.435	0.152	(6.01)	(0.00)
Volatility	0.100	0.103	-0.004	(-0.62)	(0.56)
Skewness	-0.315	-0.478	0.163	(2.87)	(0.00)
Kurtosis	5.158	5.692	-0.533	(-1.86)	(0.01)

Panel B: Fund Flows

	Main Fund	Lyxor	Difference	t-statistic	Wilcox p-value
Mean (Absolute)	0.088	0.176	-0.088	(-8.72)	(0.00)
Median (Absolute)	0.035	0.083	-0.048	(-11.15)	(0.00)
Standard Deviation	0.118	0.202	-0.084	(-7.55)	(0.00)

Table 3.4: Cross-sectional regression of difference in first order moving average coefficient between main fund and Lyxor account on fund characteristics

Results of a cross-sectional regression of the average difference in the MA(1) coefficients between the main funds and corresponding Lyxor separate accounts on main fund characteristics. The regression equation is

$$\begin{aligned}\theta_{i,dif} =& \alpha_i + \beta_1 \text{LiquidationTime}_i + \beta_2 \text{MinInvest}_i + \beta_3 \text{RedemptionFreq}_i \\ &+ \beta_4 \text{AumRatio}_i + \epsilon_i\end{aligned}$$

where $\theta_{i,dif}$ is the difference between the MA(1) coefficients of the main fund i and its associated separate account. Total liquidation time is the sum of the lockup period length and redemption notice period length for fund i, measured in years. The minimum investment is reported in millions of US dollars. The redemption frequency is the time between redemption windows, in years. AumRatio_i represents the ratio of the average AUM of separate account i to the average AUM of main fund i. The sample period is from 2002 to 2010. t-statistics are reported in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	(1)	(2)	(3)	(4)	(5)
Intercept	0.045***	0.065***	0.058***	0.046***	0.032
	(2.76)	(4.36)	(3.50)	(2.64)	(1.43)
Total Liquidation Time	0.124*				0.126*
	(1.90)				(1.91)
Minimum Investment		-0.003			-0.002
		(-0.84)			(-0.63)
Redemption Frequency			0.020		0.010
			(0.38)		(0.19)
Lyxor/Main AUM Ratio				0.025	0.025
				(1.38)	(1.37)
R^2	2.6%	0.5%	0.1%	1.4%	4.5%

Table 3.5: Cross-sectional regression of average performance difference between main fund and Lyxor account on fund characteristics

Results from a cross-sectional regression of the average performance difference between the main funds and the corresponding Lyxor separate accounts on fund characteristics. The regression equation is

$$\begin{aligned}\tilde{r}_{i,dif} =& \alpha_i + \beta_1 \text{LiquidationTime}_i + \beta_2 \text{MinInvest}_i + \beta_3 \text{RedemptionFreq}_i \\ &+ \beta_4 \text{AumRatio}_i + \epsilon_i\end{aligned}$$

where $\tilde{r}_{i,dif}$ is the difference in average monthly returns between the main fund i and its separate account. Total liquidation time is the sum of the lockup period length and redemption notice period length for fund i, measured in years. The minimum investment is reported in millions of US dollars. The redemption frequency is the time between redemption windows, in years. AumRatio_i represents the ratio of the average AUM of separate account i to the average AUM of main fund i. The sample period is from 2002 to 2010. t-statistics are reported in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	(1)	(2)	(3)	(4)	(5)
Intercept	0.119***	0.150***	0.144***	0.174***	0.162***
	(5.34)	(7.24)	(6.31)	(7.26)	(5.35)
Total Liquidation Time	0.192**				0.186**
	(2.15)				(2.08)
Minimum Investment		-0.003			-0.004
		(-0.75)			(-1.05)
Redemption Frequency			0.004		-0.019
			(0.05)		(-0.28)
Lyxor/Main AUM Ratio				-0.051**	-0.053**
				(-2.07)	(-2.13)
R^2	3.4%	0.4%	0.0%	3.1%	7.0%

Table 3.6: Cross-sectional regressions of difference in performance or MA(1) smoothing coefficient between between main fund and Lyxor account on flow measures

Panel A shows the result of a cross-sectional regression of the average performance difference between the main funds and the corresponding Lyxor separate accounts on the standard deviation of flows. Dollar flows are normalized by the time-series average AUM by fund. That is,

$$\text{NormFlow}_{i,t} = \frac{\text{AUM}_{i,t} - \text{AUM}_{i,t-1}(1 + r_{i,t})}{\frac{1}{T}\sum_{s=1}^{T} \text{AUM}_{i,s}}.$$

The dependent variable is the standard deviation of flows, normalized in this manner. Let $\text{SD}(\text{NormFlow}_i)$ be the standard deviation of normalized flows for the fund i. The regression equation is

$$\tilde{r}_{i,dif} = \alpha_i + \beta_1 \text{SD}(\text{LyxorNormFlow}_i) + \beta_2 \text{SD}(\text{MainFundNormFlow}_i) + \epsilon_i$$

where $\tilde{r}_{i,dif}$ is the difference in average monthly returns between the main fund i and its separate account. In Panel B, the dependent variable is instead the difference in MA(1) coefficient between the main fund and separate accounts. The regression equation is

$$\theta_{i,dif} = \alpha_i + \beta_1 \text{SD}(\text{LyxorNormFlow}_i) + \beta_2 \text{SD}(\text{MainFundNormFlow}_i) + \epsilon_i$$

In both cases the sample period is from 2002 to 2010. t-statistics are reported in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

Panel A: Performance differences on flow

Dep var: Difference in average returns	(1)	(2)	(3)
Intercept	0.054	0.207***	0.106*
	(1.20)	(5.32)	(1.85)
SD(Lyxor Normalized Flow)	0.447**		0.505**
	(2.22)		(2.33)
SD(Main Fund Normalized Flow)		-0.543*	-0.539*
		(-1.93)	(-1.95)
R^2	3.6%	3.0%	7.1%

Panel B: Differences in MA(1) on flow

Dep var: Difference in MA(1) coefficients	(1)	(2)	(3)
Intercept	0.027	0.047	0.003
	(0.81)	(1.64)	(0.07)
SD(Lyxor Normalized Flow)	0.169		0.220
	(1.14)		(1.36)
SD(Main Fund Normalized Flow)		0.122	0.124
		(0.59)	(0.60)
R^2	1.0%	0.3%	1.8%

Table 3.7: Regression analysis of Lyxor account flows and past performance

Results of a fund fixed-effect regression of monthly flows into Lyxor accounts on past flows and monthly returns to that account:

$$\text{Flow}_{i,t} = \sum_{j=1}^{4} (\beta_j r_{i,t-j} + \gamma_j \text{Flow}_{i,t-j}) + \epsilon_{i,t}$$

where $r_{i,t}$ is the return reported by fund i in month t; $\text{Flow}_{i,t}$ is the flow into Lyxor separate account i in month t, calculated as

$$\text{Flow}_{i,t} = \frac{\text{AUM}_{i,t} - \text{AUM}_{i,t-1}(1 + r_{i,t})}{\text{AUM}_{i,t-1}}$$

and $\text{AUM}_{i,t}$ is the AUM of Lyxor separate account i in month t. The sample period is from 2002 to 2010. Standard errors are clustered by fund and corrected for heteroskedasticity. t-statistics are reported in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

Dep var: Lyxor Flow_t	(1)	(2)	(3)	(4)
Lyxor r_{t-1}	1.658***	1.647***	1.660***	1.640***
	(13.31)	(14.00)	(14.62)	(14.40)
Lyxor r_{t-2}		0.526***	0.432***	0.428***
		(4.50)	(3.75)	(3.68)
Lyxor r_{t-3}			0.274**	0.307***
			(2.51)	(2.77)
Lyxor r_{t-4}				-0.076
				(-0.60)
Lyxor Flow_{t-1}	0.084***	0.102***	0.090***	0.085***
	(10.56)	(8.06)	(5.97)	(5.02)
Lyxor Flow_{t-2}		0.024***	0.032***	0.026**
		(3.68)	(2.97)	(2.01)
Lyxor Flow_{t-3}			0.021***	0.031***
			(3.16)	(2.95)
Lyxor Flow_{t-4}				0.016***
				(2.67)
R^2	4.3%	6.0%	6.6%	6.5%

Table 3.8: Regression analysis of main fund flows and past performance

Results of a fund fixed-effect regression of monthly flows into main hedge funds on past flows and monthly returns to that account

$$\text{Flow}_{i,t} = \sum_{j=1}^{4} (\beta_j r_{i,t-j} + \gamma_j \text{Flow}_{i,t-j}) + \epsilon_{i,t}$$

where $r_{i,t}$ is the return reported by fund i in month t; $\text{Flow}_{i,t}$ is the flow into main hedge fund i in month t, calculated by

$$\text{Flow}_{i,t} = \frac{\text{AUM}_{i,t} - \text{AUM}_{i,t-1}(1 + r_{i,t})}{\text{AUM}_{i,t-1}},$$

and $\text{AUM}_{i,t}$ is the AUM of main hedge fund i in month t. The sample period is from 2002 to 2010. Standard errors are clustered by fund and corrected for heteroskedasticity. t-statistics are reported in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

Dep var: Main Fund Flow_t	(1)	(2)	(3)	(4)
Main Fund r_{t-1}	0.347	0.279	0.297	0.286
	(1.29)	(1.00)	(1.06)	(0.99)
Main Fund r_{t-2}		0.574***	0.493**	0.504**
		(2.94)	(2.52)	(2.47)
Main Fund r_{t-3}			0.397	0.340
			(1.10)	(0.94)
Main Fund r_{t-4}				0.265
				(1.14)
Main Fund Flow_{t-1}	0.007	0.004	0.003	0.002
	(0.52)	(0.38)	(0.24)	(0.15)
Main Fund Flow_{t-2}		-0.009**	-0.010***	-0.011***
		(-2.56)	(-3.02)	(-3.62)
Main Fund Flow_{t-3}			-0.006*	-0.007**
			(-1.70)	(-2.05)
Main Fund Flow_{t-4}				-0.005**
				(-2.21)
R^2	0.0%	0.1%	0.1%	0.1%

Table 3.9: Regression analysis of main fund returns and past separate account performance

Results of a fund fixed-effect regression of monthly returns in hedge funds main funds on past flows and monthly returns on the associated Lyxor separate account. The regression equation is

$$\text{Main Fund } r_{i,t} = \sum_{j=1}^{3} \left(\beta_j \text{Lyxor } r_{i,t-j} + \gamma_j \text{Lyxor Flow}_{i,t-j}\right) + \epsilon_{i,t}$$

where Main Fund $r_{i,t}$ and Lyxor $r_{i,t}$ are the returns to the main fund and associated separate account for fund i in month t and Lyxor Flow$_{i,t}$ are the flows associated with the separate account for fund i in month t. The sample period is from 2002 to 2010. Standard errors are clustered by fund and corrected for heteroskedasticity. t-statistics are reported in parentheses. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

Dep var: Main Fund $r_{i,t}$	(1)	(2)	(3)	(4)
Lyxor r_{t-1}	0.145***	0.152***	0.150***	0.149***
	(9.53)	(9.81)	(9.60)	(9.45)
Lyxor r_{t-2}		-0.030**	-0.028**	-0.024*
		(-2.17)	(-1.99)	(-1.71)
Lyxor r_{t-3}			-0.032**	-0.036**
			(-2.22)	(-2.39)
Lyxor r_{t-4}				0.063***
				(4.49)
Lyxor Flow$_{t-1}$	0.000	0.000	0.000	0.000
	(-0.13)	(-0.32)	(0.08)	(0.13)
Lyxor Flow$_{t-2}$		0.001*	0.002*	0.001
		(1.76)	(1.88)	(0.41)
Lyxor Flow$_{t-3}$			0.000	0.000
			(0.50)	(-0.35)
Lyxor Flow$_{t-4}$				0.000
				(-0.26)
R^2	2.2%	2.3%	2.4%	2.7%

Vita

Grant Verdell Farnsworth

Education

Pennsylvania State University, University Park, PA

Ph.D. Finance, 2015 (expected)

Northwestern University, Evanston, IL

M.S. Finance, 2008

Brigham Young University, Provo, UT

B.S. Mathematics and B.A. Economics, 2004

Interests

Hedge Funds, Mutual Funds, Financial Institutions, Empirical Asset Pricing

Research

"Essays on the Structure, Performance, and Behavior of Hedge Funds", *Job Market Paper*. Presentations: Penn State University (2015), University of Arizona (2015), Syracuse University (2015), Texas Christian University (2015), Federal Reserve Board of Governors (2015), University of Kansas (2015)

"Liquidity Costs, Return Smoothing, and Investor Flows: Evidence from a Separate Account Platform", with Charles Cao, Bing Liang, and Andrew Lo. Presentations: 2013 SFS Finance Cavalcade, 5th Annual Hedge Fund Conference (2013), Midwest Finance Conference (2013), Penn State University (2013)

"Dollars and Sense: Investor Demand, Managerial Skill, and Hedge Fund Startups", with Charles Cao and Hong Zhang. Presentations: 2015 SFS Finance Cavalcade