

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

Charles Cao,^a Grant Farnsworth,^{a,b} Bing Liang,^c Andrew W. Lo^d

^aSmeal College of Business, Pennsylvania State University, University Park, Pennsylvania 16802; ^bNeeley School of Business, Texas Christian University, Fort Worth, Texas 76109; ^cIsenberg School of Management, University of Massachusetts Amherst, Amherst, Massachusetts 01103; ^dSloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142

Contact: qxc2@psu.edu (CC); g.v.farnsworth@tcu.edu (GF); bliang@isenberg.umass.edu (BL); alo-admin@mit.edu (AWL)

Received: November 24, 2014
Revised: May 30, 2015
Accepted: August 9, 2015
Published Online in Articles in Advance: April 29, 2016

<https://doi.org/10.1287/mnsc.2015.2401>

Copyright: © 2016 INFORMS

Abstract. We use a new hedge fund data set 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 and (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 a 1.7% performance reduction on an annual basis. Investor flows chase monthly past performance on the platform but not in the associated funds.

History: Accepted by Neng Wang, finance.

Keywords: hedge funds • separate accounts • return smoothing • share restrictions

1. Introduction

Hedge funds returns are conspicuous among investment vehicles because their voluntarily reported (typically nonaudited) 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 suffer significantly?

As 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 we typically do not observe the trades or mark-to-market methods, so we 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 to examine 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 as in Agarwal et al. (2011) and Bollen and Pool (2009).² 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 et al. (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 as a result of flow-induced trading, and cash management.

This paper is also unique in that it directly investigates the relation between share restrictions 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 (2014) examine the performance difference between Undertakings for Collective Investment in Transferable Securities (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.³ 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 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.

From a methodological perspective, we construct a merged main fund and separate account data set 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 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, whereas the associated Lyxor accounts have, on average, a statistically different coefficient of 0.121. Since underlying asset 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 as a result of managerial discretion is higher for funds with greater barriers to liquidity as measured by the total time to redemption 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, although reducing 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, whereas 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 2 develops our hypotheses of interest and places them in the context of the existing literature. Section 3 describes the Lyxor separate account platform, our data set, and the methodology used to construct a sample appropriate for testing. Section 4 then outlines the econometric tests used and discusses the results. Section 5 concludes the paper.

2. Hypothesis Development

2.1. Third-Party Valuation, Fund Liquidity, and Return Smoothing

Self-reported hedge fund returns are serially correlated (see Asness et al. 2001, Getmansky et al. 2004, Loudon et al. 2006). That these returns are voluntarily reported and independent valuations 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 (NAVs). For example, Bollen and Pool (2008) discuss 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. Furthermore, the discontinuity disappears when using bimonthly returns, indicating that the reported returns are subsequently reversed. Agarwal et al. (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 redemptions 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 et al. (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 nonsynchronous 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 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 that results from stale prices and nonsynchronous trading will 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 2. *If share restrictions facilitate managerial discretion, then funds with more stringent share restrictions will have greater smoothing than those with fewer.*

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

Furthermore, 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 4. *If share restrictions are put in place in the main fund 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.*

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 et al. (1993).

In empirical tests of hot hands performance persistence among hedge funds, Jagannathan et al. (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 et al. (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 et al. (2014) 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. By 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.

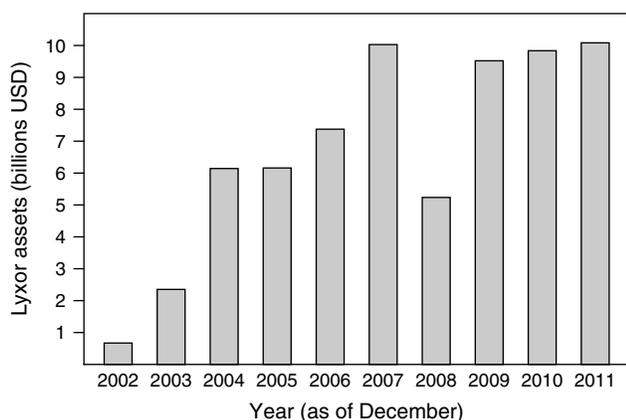
Hypothesis 5. *If hedge fund investors seek to chase fund performance, we expect a stronger positive relation between fund performance and subsequent flows for Lyxor funds than for main funds, which provide fewer opportunities for flows.*

3. Data Description

3.1. The Lyxor Separate Account Platform

We utilize a unique data set 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

Figure 1. Lyxor Assets Under Management



Notes. 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 U.S. dollars.

in assets under management (AUM). Figure 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 Deutsche Bank, Man Group, AlphaMetrix, Goldman Sachs, UBS, and other institutions. Some institutional details vary from one separate account platform to another, but 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 are 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.⁴ 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 these data, the difference between main fund and separate account performance we report is a lower bound on what investors would face.

Whereas main funds often provide self-reported performance information relatively infrequently (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, diversification difficulty, 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.

3.2. Separate Account and Main Fund Data

Lyxor publishes performance for all funds on the platform each week. Whereas 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 data set containing defunct fund performance as well. Survivorship and other biases associated with hedge fund returns, discussed in Fung and Hsieh (2000), Liang (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 data set contains weekly historical returns for 291 Lyxor accounts between 2002 and 2010. After removing funds with fewer 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 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 data sets: TASS, HFR, CISDM, BarclayHedge, and Morningstar. These data sets 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 data sets representing 135 fully vetted fund matches.

Table 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

Table 1. Summary Statistics of Hedge Fund Characteristics

	Main fund		Lyxor account		HF universe	
	Mean	Median	Mean	Median	Mean	Median
<i>Incentive fee (%)</i>	19.20	20.00	20.35	20.00	18.28	20.00
<i>Management fee (%)</i>	1.52	1.50	1.58	1.50	1.45	1.50
<i>Lockup period (months)</i>	0.57	0.00	0.00	0.00	3.49	0.00
<i>Redemption frequency (months)</i>	2.00	1.00	0.23	0.23	2.36	1.00
<i>Redemption notice period (months)</i>	1.00	1.00	0.07	0.07	1.08	1.00
<i>Minimum allowed investment (\$ millions)</i>	1.71	1.00	0.00	0.00	0.89	0.50
<i>High watermark (yes = 1)</i>	0.85	—	—	—	0.79	—
<i>Leverage allowed (yes = 1)</i>	0.61	—	—	—	0.52	—
<i>Average AUM (\$ millions)</i>	699.18	138.59	65.67	48.95	427.61	86.48

Notes. Summary statistics of the fund-level characteristics of the main funds and Lyxor separate accounts. Each Lyxor/main fund pair is an observation. *Incentive fee* and *management fee* 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. *Average AUM* is the average reported monthly AUM, expressed in millions of U.S. 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 (HF) universe during this time period are included for comparison.

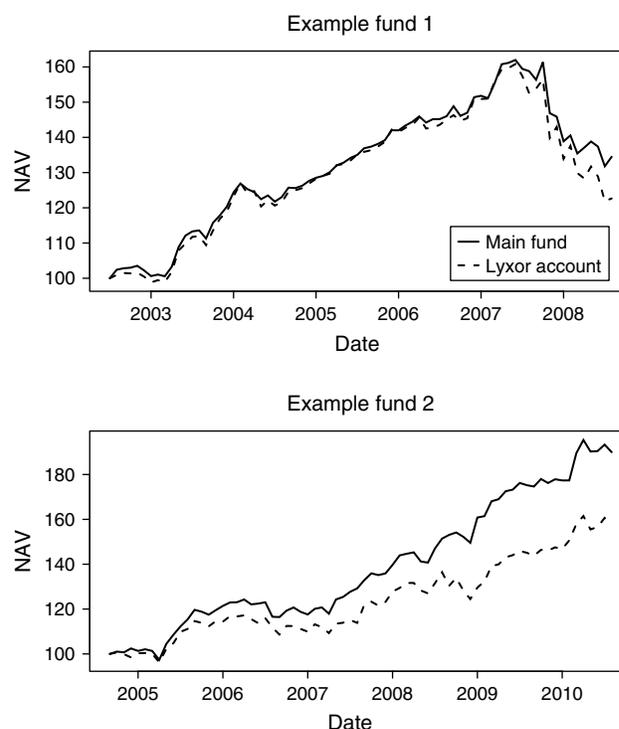
incentive fee and approximately 1.5% for the management fee. There are no lockup periods or fund-wise minimum investment amounts for the separate accounts, although a minority of main funds in our sample do report initial lockup periods (most frequently of a year) and all have minimum investment amounts, often about \$1 million. 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 two 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 or 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 data set. Lyxor accounts are typically significantly

smaller than their main fund equivalent, with the mean account AUM of approximately \$66 million (versus the mean of \$699 million in the associated main funds).

Figure 2 shows the Lyxor and main fund performance of two example funds. In the first case, the Lyxor

Figure 2. Net Asset Value of Two Representative Funds: Main Funds vs. Separate Accounts



Notes. 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; the second has more. Initial NAVs are set to 100 so that the comparison is straightforward.

account return closely matches that of the associated main fund; in the second case, there is a greater performance difference. These examples not only 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 the 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%.

4. Results

4.1. Fund-Level Tests

Getmansky et al. (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 2. We check moving average models of up to order 3, denoted as MA(1), MA(2), and MA(3) throughout.

In panel A of Table 2, we see that Lyxor funds exhibit an average MA(1) coefficient of 0.121, whereas 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 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 2, we see the coefficients from an MA(2) model. The first moving average coefficient for both the main fund and the 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

Table 2. Lyxor and Main Fund Smoothing

	Main fund	Lyxor	Difference
Panel A: MA(1) model			
Average θ_1	0.182	0.121	0.061
t -Statistic	(9.22)	(6.28)	(4.35)
Wilcoxon p -value	—	—	(0.00)
Panel B: MA(2) model			
Average θ_1	0.195	0.122	0.073
t -Statistic	(8.62)	(5.88)	(4.16)
Wilcoxon p -value	—	—	(0.00)
Average θ_2	0.025	-0.011	0.037
t -Statistic	(1.28)	(-0.54)	(2.04)
Wilcoxon p -value	—	—	(0.01)
Panel C: MA(3) model			
Average θ_1	0.197	0.112	0.085
t -Statistic	(8.77)	(5.12)	(5.01)
Wilcoxon p -value	—	—	(0.00)
Average θ_2	0.021	-0.005	0.026
t -Statistic	(0.99)	(-0.23)	(1.31)
Wilcoxon p -value	—	—	(0.01)
Average θ_3	-0.018	-0.019	0.001
t -Statistic	(-0.84)	(-0.81)	(0.07)
Wilcoxon p -value	—	—	(0.86)

Notes. The moving average process coefficients, fund by fund, for each fund and each separate account are estimated for models up to order 3:

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

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 Wilcoxon tests for differences between main fund and separate account coefficient ranks.

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

$$\frac{\gamma}{\delta + \gamma} = \frac{\theta_1^m - \theta_1^s}{\theta_1^m}. \quad (1)$$

Using the estimated coefficients in panel A of Table 2, we conclude that approximately 33% ($= 1 - 0.121/0.182$) of hedge fund smoothing appears to be attributable to managerial reporting, whereas 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 et al. (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 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: 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 Wilcoxon 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

Table 3. Distributional Properties of Lyxor and Main Fund Returns and Flows

	Main fund	Lyxor	Difference	<i>t</i> -Statistic	Wilcoxon <i>p</i> -value
Panel A: Fund returns					
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					
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)

Notes. 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, and the standard deviation is computed on signed flows. The *t*- and Wilcoxon tests are tests performed on the cross section of each fund statistic.

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. Although 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.

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 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 NAVs) 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:

$$\theta_{i,dif} = \alpha_i + \beta_1 RedemptionTime_i + \beta_2 MinInvest_i + \beta_3 RedemptionFreq_i + \beta_4 AUMRatio_i + \epsilon_i, \quad (2)$$

where $\theta_{i,dif}$ is the difference between the MA(1) coefficients estimated for fund *i* in the main fund and in

Table 4. Cross-Sectional Regression of Difference in First-Order Moving Average Coefficient Between Main Fund and Lyxor Account on Fund Characteristics

	(1)	(2)	(3)	(4)	(5)
<i>Intercept</i>	0.045*** (2.76)	0.065*** (4.36)	0.058*** (3.50)	0.046*** (2.64)	0.032 (1.43)
<i>Total redemption time</i>	0.124* (1.90)				0.126* (1.91)
<i>Minimum investment</i>		−0.003 (−0.84)			−0.002 (−0.63)
<i>Redemption frequency</i>			0.020 (0.38)		0.010 (0.19)
<i>Lyxor/Main AUM ratio</i>				0.025 (1.38)	0.025 (1.37)
<i>R</i> ² (%)	2.6	0.5	0.1	1.4	4.5

Notes. 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

$$\theta_{i,\text{dif}} = \alpha_i + \beta_1 \text{RedemptionTime}_i + \beta_2 \text{MinInvest}_i + \beta_3 \text{RedemptionFreq}_i + \beta_4 \text{AUMRatio}_i + \epsilon_i,$$

where $\theta_{i,\text{dif}}$ is the difference between the MA(1) coefficients of the main fund i and its associated separate account. *Total redemption 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 U.S. dollars. The *redemption frequency* is the time between redemption windows, in years. *Lyxor/Main AUM ratio* 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.

***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

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 redemption for main fund i , RedemptionTime_i , in years. The minimum permitted investment in main fund i is denoted by MinInvest_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 the 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 4 reports the results of these regressions. We find that the total redemption time, a measure of fund-level illiquidity, is positively and significantly associated with greater moving average smoothing coefficients in the main fund compared with 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 redemption 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. Although 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 4 that funds with higher share restrictions have a greater share of "managerial" smoothing supports Hypothesis 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 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.

Table 5. Cross-Sectional Regression of Average Performance Difference Between the Main Fund and Lyxor Accounts on Fund Characteristics

	(1)	(2)	(3)	(4)	(5)
<i>Intercept</i>	0.119*** (5.34)	0.150*** (7.24)	0.144*** (6.31)	0.174*** (7.26)	0.162*** (5.35)
<i>Total redemption time</i>	0.192** (2.15)				0.186** (2.08)
<i>Minimum investment</i>		-0.003 (-0.75)			-0.004 (-1.05)
<i>Redemption frequency</i>			0.004 (0.05)		-0.019 (-0.28)
<i>Lyxor/Main AUM ratio</i>				-0.051** (-2.07)	-0.053** (-2.13)
<i>R² (%)</i>	3.4	0.4	0.0	3.1	7.0

Notes. 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

$$\bar{r}_{i,dif} = \alpha_i + \beta_1 \text{RedemptionTime}_i + \beta_2 \text{MinInvest}_i + \beta_3 \text{RedemptionFreq}_i + \beta_4 \text{AUMRatio}_i + \epsilon_i,$$

where $\bar{r}_{i,dif}$ is the difference in average monthly returns between the main fund i and its separate account. *Total redemption 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 U.S. dollars. The *redemption frequency* is the time between redemption windows, in years, and *Lyxor/Main AUM ratio* 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.

***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

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 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:

$$\bar{r}_{i,dif} = \text{MainFund}\bar{r}_i - \text{Lyxor}\bar{r}_i, \quad (3)$$

where $\text{MainFund}\bar{r}_i$ denotes the time-series average return over the sample periods to the main fund i , and $\text{Lyxor}\bar{r}_i$ denotes the corresponding time-series average return to that fund’s associated Lyxor separate account. Table 5 reports the results from the following cross-sectional regression:

$$\bar{r}_{i,dif} = \alpha_i + \beta_1 \text{RedemptionTime}_i + \beta_2 \text{MinInvest}_i + \beta_3 \text{RedemptionFreq}_i + \beta_4 \text{AUMRatio}_i + \epsilon_i. \quad (4)$$

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 redemption 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. Recall 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.

As seen in Table 5, total redemption 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 total redemption time implies a monthly performance difference between the main fund and separate account of 0.186 (t -statistic = 2.08). Both components of total redemption 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 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 redemption time faced by a new investor, have a greater serial correlation difference between the 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 as a result of 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.

4.3. The Cross Section of Flow Volatility

The fund characteristics used in Section 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 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,

$$NormFlow_{i,t} = \frac{AUM_{i,t} - AUM_{i,t-1}(1 + r_{i,t})}{(1/T) \sum_{s=0}^T AUM_{i,s}}. \quad (5)$$

This measure looks both forward and backward in the time series 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 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 (coefficient = 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.

Table 6. Cross-Sectional Regressions of Difference in Performance or MA(1) Smoothing Coefficient Between the Main Fund and Lyxor Accounts on Flow Measures

Panel A: Performance differences on flow			
DV: Difference in average returns	(1)	(2)	(3)
<i>Intercept</i>	0.054 (1.20)	0.207*** (5.32)	0.106* (1.85)
<i>SD(Lyxor Normalized Flow)</i>	0.447** (2.22)		0.505** (2.33)
<i>SD(Main Fund Normalized Flow)</i>		-0.543* (-1.93)	-0.539* (-1.95)
<i>R</i> ² (%)	3.6	3.0	7.1

Panel B: Differences in MA(1) on flow			
DV: Difference in MA(1) coefficients	(1)	(2)	(3)
<i>Intercept</i>	0.027 (0.81)	0.047 (1.64)	0.003 (0.07)
<i>SD(Lyxor Normalized Flow)</i>	0.169 (1.14)		0.220 (1.36)
<i>SD(Main Fund Normalized Flow)</i>		0.122 (0.59)	0.124 (0.60)
<i>R</i> ² (%)	1.0	0.3	1.8

Notes. 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 (SD) of flows. Dollar flows are normalized by the time-series average AUM by fund. That is,

$$NormFlow_{i,t} = \frac{AUM_{i,t} - AUM_{i,t-1}(1 + r_{i,t})}{(1/T) \sum_{s=1}^T AUM_{i,s}}$$

The dependent variable (DV) is the standard deviation of flows, normalized in this manner. Let $SD(NormFlow_i)$ be the standard deviation of normalized flows for the fund i . The regression equation is

$$\bar{r}_{i,dif} = \alpha_i + \beta_1 SD(LyxorNormFlow_i) + \beta_2 SD(MainFundNormFlow_i) + \epsilon_i,$$

where $\bar{r}_{i,dif}$ 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 the MA(1) coefficient between the main fund and separate accounts. The regression equation is

$$\theta_{i,dif} = \alpha_i + \beta_1 SD(LyxorNormFlow_i) + \beta_2 SD(MainFundNormFlow_i) + \epsilon_i.$$

In both cases, the sample period is from 2002 to 2010. t -Statistics are reported in parentheses.

***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

In panel B of Table 6 we report the results of a similar regression in which we use the standard deviation of flows to explain differences in the 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 the MA(1) coefficient between main funds and separate accounts is manager discretion in how to calculate and report returns.

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

$$Flow_t = \frac{AUM_t - AUM_{t-1}(1 + r_t)}{AUM_{t-1}}, \quad (6)$$

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-effect regression:

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

where $Flow_{i,t}$ 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 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 are 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 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 8 reports these results and shows a different picture. Flow autocorrelation is not significant in the first lag and is negative

Table 7. Regression Analysis of Lyxor Account Flows and Past Performance

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

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

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

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

$$Flow_{i,t} = \frac{AUM_{i,t} - AUM_{i,t-1}(1 + r_{i,t})}{AUM_{i,t-1}};$$

and $AUM_{i,t}$ 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. DV, dependent variable.

***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

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)–(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 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.

4.5. Main Fund Return Prediction

Since separate account returns are calculated by a third party and reported almost immediately, whereas main fund returns are calculated by the fund itself and reported relatively infrequently, there may be investors who have access to updated Lyxor performance

Table 8. Regression Analysis of Main Fund Flows and Past Performance

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

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

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

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

$$Flow_{i,t} = \frac{AUM_{i,t} - AUM_{i,t-1}(1 + r_{i,t})}{AUM_{i,t-1}};$$

and $AUM_{i,t}$ 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. DV, dependent variable.

***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

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 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

$$Main\ Fund\ r_{i,t} = \sum_{j=1}^4 (\beta_j Lyxor\ r_{i,t-j} + \gamma_j Lyxor\ Flow_{i,t-j}) + \epsilon_{i,t}, \quad (8)$$

where $Main\ Fund\ r_{i,t}$ is the return in main fund i in month t and $Lyxor\ r_{i,t}$ is the associated separate

Table 9. Regression Analysis of Main Fund Returns and Past Separate Account Performance

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

Notes. 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

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

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 , respectively, 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. DV, dependent variable.

***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

account return in month t . Additionally, Lyxor 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.

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

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

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 eight funds (5.9%), β_2 is negative and significant, and for three 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

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

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, whereas 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.

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, in either the short or the long run. Here, we discuss some possibilities.

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, whereas 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.⁵

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 U.S. 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 is the cost of fund-level liquidity.

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

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, nonmimicking 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 6, 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 Pástor and Stambaugh (2003) liquidity factor with controls for the market excess return, small minus big (*SMB*) size factor, high minus low (*HML*) value/growth factor, and momentum (*MOM*) factor. The regression equation is

$$r_{i,t} = \alpha_i + \gamma_1 LIQ_t + \gamma_2 MktRf_t + \gamma_3 SMB_t + \gamma_4 HML_t + \gamma_5 MOM_t + \epsilon_{i,t}. \quad (11)$$

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 what 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. Although 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.

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 the performance reduction resulting from eliminating hedge fund share restrictions is.

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 redemption, appear to have greater disparity in the 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 redemption. 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.

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, whereas 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.

Acknowledgments

The authors gratefully acknowledge the insights and comments of Jayna Cummings, Serge Darolles, Heber Farnsworth, Will Goetzmann, Brad Goldie, Jeremiah Green, Craig Holden, David Hsieh, Ravi Jagannathan, Stefan Keller, Olga Kolokolova, Bill Krakow, Pete Kyle, Lionel Paquin, Luboš Pástor, Veronika Pool, Ronnie Sadka, Tim Simin, Jess Smith, and Melvyn Teo; and seminar/conference participants at Pennsylvania State University, the Fifth Annual Hedge Fund Research Conference in Paris, and the 2013 SFS Finance Cavalcade in Miami. The authors are especially grateful to Lyxor Asset Management, a subsidiary of Société Générale, for providing the separate account performance data used in this paper. All remaining errors are those of the authors. The views and opinions expressed in this article are those of the authors only and do not necessarily represent the views and opinions of any institution or agency, any of their affiliates or employees, or any of the individuals acknowledged above.

Endnotes

¹ From the point of view of an investor, cash invested in a hedge fund may be impossible to liquidate rapidly because of 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 whereas 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.

² It should be noted that Jorion and Schwarz (2013) suggest that the return discontinuity noted by Bollen and Pool (2009) may be explained by the incentive fee structure.

³UCITS funds conform to a set of European Union (EU) directives that allow funds to operate freely in any EU state on the basis of authorization from a single state.

⁴Although 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.

⁵Thanks to Lionel Paquin and Stefan Keller from Lyxor Asset Management, who provided this insight about Lyxor's separate accounts.

References

- Agarwal V, Naik N (2000) Multi-period performance persistence analysis of hedge funds. *J. Financial Quant. Anal.* 35(3):327–342.
- Agarwal V, Daniel N, Naik N (2004) Flows, performance, and managerial incentives in hedge funds. *EFA 2003 Annual Conf.*, paper no. 501, European Finance Association, Brussels. <http://ssrn.com/abstract=424369>.
- Agarwal V, Daniel N, Naik N (2011) Do hedge funds manage their reported returns? *Rev. Financial Stud.* 24(10):2221–2256.
- Ang A, Bollen N (2010) Locked up by a lockup: Valuing liquidity as a real option. *Financial Management* 39(3):1069–1096.
- Aragon G (2007) Share restrictions and asset pricing: Evidence from the hedge fund industry. *J. Financial Econom.* 83(1):33–58.
- Aragon G, Liang B, Park H (2014) Onshore and offshore hedge funds: Are they twins? *Management Sci.* 60(1):74–91.
- Asness C, Kraill R, Liew J (2001) Do hedge funds hedge? *J. Portfolio Management* 28(1):6–19.
- Beck N, Katz J (1995) What to do (and not to do) with time-series cross-section data. *Amer. Political Sci. Rev.* 89(3):634–647.
- Bollen N, Pool V (2008) Conditional return smoothing in the hedge fund industry. *J. Financial Quant. Anal.* 43(2):267–298.
- Bollen N, Pool V (2009) Do hedge fund managers misreport returns? evidence from the pooled distribution. *J. Finance* 64(5):2257–2288.
- Fung W, Hsieh D (2000) Performance characteristics of hedge funds and commodity funds: Natural vs. spurious biases. *J. Financial Quant. Anal.* 35(3):291–307.
- Getmansky M, Lo A, Makarov I (2004) An econometric model of serial correlation and illiquidity in hedge fund returns. *J. Financial Econom.* 74(3):529–609.
- Hendricks D, Patel J, Zeckhauser R (1993) Hot hands in mutual funds: Short-run persistence of relative performance, 1974–1988. *J. Finance* 48(1):93–130.
- Jagannathan R, Malakhov A, Novikov D (2010) Do hot hands exist among hedge fund managers? An empirical evaluation. *J. Finance* 65(1):217–255.
- Joenvaara J, Kosowski R (2014) Geography, liquidity and fund performance: New evidence from UCITS hedge funds. Working paper, Imperial College Business School, London.
- Jorion P, Schwarz C (2013) The delisting bias in hedge fund databases. *J. Alternative Investments* 16(3):37–47.
- Liang B (1999) On the performance of hedge funds. *Financial Analysts J.* 55(4):72–85.
- Liang B (2000) Hedge funds: The living and the dead. *J. Financial Quant. Anal.* 35(3):309–326.
- Loudon G, Okunev J, White D (2006) Hedge fund risk factors and the value at risk of fixed income trading strategies. *J. Fixed Income* 16(2):46–61.
- Ozik G, Sadka R (2015) Skin in the game versus skimming the game: Governance, share restrictions, and insider flows. *J. Financial Quant. Anal.* 50(6):1293–1319.
- Pástor L, Stambaugh RF (2003) Liquidity risk and expected stock returns. *J. Political Econom.* 111(3):642–685.
- Sirri E, Tufano P (1998) Costly search and mutual fund flows. *J. Finance* 53(5):1589–1622.
- Zheng L (1999) Is money smart? A study of mutual fund investors' fund selection ability. *J. Finance* 54(3):901–933.