Illiquidity in Intermediary Portfolios: Evidence from Large Hedge Funds[1]

Daniel Barth,[2,3,4] Phillip Monin[5]

June 16, 2018

Abstract

This paper uses the first systematic, regulatory collection of data on large hedge funds to address three empirical questions: how large is the return premium earned on illiquid assets, how important is the illiquidity premium for explaining hedge funds' risk-adjusted returns, and who ultimately captures the illiquidity premium, hedge fund managers or fund investors? We argue that the potential impact from forced asset sales comprises a compensated risk for hedge fund managers. In response, managers restrict fund shares and charge higher fees. We find that investor share restrictions and portfolio illiquidity are tightly linked, suggesting managers pass much of the costs of illiquid assets to fund investors. We estimate an illiquidity premium of 4.3% per year between the 5th and 95th percentiles of fund illiquidity, equivalent to an additional 82 basis points per year for an additional log-day of illiquidity. We also find that 38% of hedge fund alpha can be explained by portfolio illiquidity, but that investor share illiquidity can explain 64% of alpha. Further, coefficient estimates from an asset pricing model suggest that fund investors capture roughly 75% of the illiquidity premium. Finally, consistent with managers requiring additional compensation for undiversifiable illiquidity risk, we show that managers of illiquid funds charge higher incentive fees both unconditionally and after controlling for manager skill and a host of other controls.

JEL Classifications: G11, G12, G23

Keywords: Hedge Funds, Liquidity, Manager Skill, Investors

---

[1]This paper was previously circulated under the title "Skill, Illiquidity Risk, and Investor Preferences in Large Hedge Funds"

[2]Views and opinions expressed are those of the authors and do not necessarily represent official positions or policy of the OFR or the U.S. Department of the Treasury.

[3]We thank George Aragon, Greg Brown, Greg Duffee, Christian Lundblad, Stephen Shore, Emil Siriwardane, Russ Wermers, and seminar participants at the Federal Reserve Board of Governors, the New York Federal Reserve, American University, and the Institute for Private Capital Spring Symposium at UNC Chapel Hill. All remaining errors are our own.

[4]Office of Financial Research, U.S. Department of Treasury, 717 14th St. NW, Washington, DC, 20005, USA. Phone: (202) 927-8235. Email: Daniel.Barth@ofr.treasury.gov

[5]Office of Financial Research, U.S. Department of Treasury, 717 14th St. NW, Washington, DC, 20005, USA. Phone: (202) 927-8277. Email: Phillip.Monin@ofr.treasury.gov

1 Introduction

(4) Do plots for each strategy separately

(6) Cut by median liquidity to show premium is similar for subgroups

Illiquidity is a well-established source of heterogeneity in the cross-section of asset returns. Unfortunately, measuring the return premium associated with illiquid assets is a difficult undertaking. Nearly all measures of illiquidity result from data based on trades, such as price impact per volume or short-term reversals.[6] Yet, assets that trade sufficiently often to produce credible estimates of such measures are by construction the most liquid. Measuring the illiquidity premium across the full spectrum of illiquidity, including deeply illiquid assets that trade infrequently or in some cases not at all, is impossible from the trade-based measures of illiquidity available in the literature.[7]

In this paper, we estimate the illiquidity premium through the returns to portfolios of illiquid assets held by traders who are likely to be the marginal investors in illiquid markets: hedge funds. Our measures of illiquidity — the number of days needed to liquidate assets without price impact, and the number of days needed to meet investor redemptions — span liquidation and redemption horizons up to 366 days, are broad enough to apply across multiple strategies and asset classes, and are not fully explained by existing measures such as return smoothing, lockup periods, or redemption notice periods.

Specifically, we address three empirical questions related illiquidity: how large is the illiquidity premium, how important is it for explaining hedge funds' risk-adjusted returns, and who ultimately captures the illiquidity premium, the fund manager or fund investors?[8] Our analysis focuses on hedge funds because hedge funds are likely to be important participants in illiquid markets. Other market participants, such as public mutual funds or pensions, are restricted in their ability to invest significant capital in illiquid assets[9], and the implicit leverage associated with outside capital is likely necessary to justify the costly expertise needed to specialize in such markets. The implication is that privately held, managed funds such as hedge funds are likely to be the specialists in illiquid asset markets. Consistent with this hypothesis, we document that many funds have substantial investments in deeply illiquid assets.

We begin with an economic argument for who ultimately bears the costs of illiquid assets, the hedge fund manager or hedge fund investors, and therefore who ultimately captures the illiquidity premium. We

[6]See Amihud and Mendelson (1986), Acharya and Pedersen (2005), Pastor and Stambaugh (2003), and Sadka (2006) as prominent examples.

[7]This is the motivation given in Khandani and Lo (2011) for using return autocorrelations to estimate illiquidity premia.

[8]Strictly speaking, the fund managers are also fund investors, but we reserve the terminology *investors* for outside investors who do not participate in the fund's investment decisions.

[9]Such restrictions may come directly from the Securities and Exchange Commission, or from internal risk management practices.

argue that the level of illiquidity of hedge fund assets constitutes an undiversifiable risk for hedge fund managers that will require compensation in equilibrium. This argument is based on theories of limits to arbitrage that predict risks arising from trading in specialized markets may be compensated, even if such risks are in principle diversifiable.[10] In markets that comprise a limited number of traders and require costly and idiosyncratic knowledge to price securities, the marginal investor is likely to be a specialist rather than a well-diversified trader. In these markets, quickly finding a buyer who is willing or able to pay fair value for the asset may not be possible. Immediate or unforeseen cash demands may force traders to sell assets at a discount, or in the case of severe market distress, at fire sale prices. Empirical support for this mechanism is prevalent; previous studies have documented large and persistent price impacts due to forced sales by institutional traders, including asset managers (Coval and Stafford (2007), Jotikasthira, Lunblad, and Ramadorai (2012), Falato, Hortascu, Li, and Shin (2017)).

The possibility that the hedge fund manager may not be able to finance a position throughout the life of the trade is referred to as *funding liquidity risk*. Hedge funds get funding from two sources: outside investor capital and borrowing. Forced asset sales can therefore originate either from investor redemptions or margin calls on leveraged positions. Consider first the investor redemption channel. Because of hedge fund managers' substantial investment of personal wealth in the fund[11], the manager (along with the remaining investors) are the "residual claimants" on the fund's assets after outside investors have redeemed their capital. If investor redemptions require assets to be sold at a discount, it is the manager and the remaining investors in the fund who will be adversely affected. Because of fund managers' large personal stakes in their funds , they will be unable to fully diversify this risk. Faced with the undiversifiable risk of price impact from investor redemptions, managers will respond in one of two ways: either by requiring explicit compensation in the form of higher fees, or through the imposition of investor share restrictions (or both).

This argument generates a number of empirical predictions. If managers retain the entirety of the illiquidity risk and charge higher fees, illiquidity premium will been found in gross-of-fee but not net-of-fee returns, and funds' asset illiquidity will be positively correlated with fees but not with investor share restrictions. Alternatively, if the manager imposes sufficiently tight share restrictions to ensure investor redemptions will never require assets to be sold too quickly, investors should capture the entire illiquidity premium and the illiquidity premium estimated from gross returns should be identical to that estimated from net returns. Further, asset illiquidity will be positively correlated with share restrictions but *negatively* associated with fees. The negative association with fees results because fees are charged on total returns; if investors

[10]See Shleifer and Vishny (1997), Gromb and Vayanos (2002), Kyle and Xiong (2001), He and Krishnamurthy (2012), Gabaix, Krishnamurthy, and Vigneron (2007), and Siriwardane (2018) as examples.

[11]This is a common feature of hedge funds intended to minimize principal-agent concerns.

capture all of the illiquidity premium, a fund who's returns originate in part from illiquidity must charge lower fees than a fund with similar performance that holds only liquid assets. Finally, if managers and investors share the illiquidity premium, or if forced asset sales may also arise from margin calls on leveraged positions, we should see illiquidity premium in both net and gross returns, and the difference in the illiquidity premium estimated from gross and net returns will estimate the proportion of the illiquidity premium passed down to fund investors. Moreover, both fees and share restrictions will be positively correlated with asset illiquidity.

Based on this argument, our first set of results establish the empirical relationship between *portfolio illiquidity*, measured as the weighted-average number of days needed to sell assets without price impact, and *investor share illiquidity*, measured as the weighted-average number of days until investors' capital is returned once a redemption is requested, each of which is reported by the funds themselves. Our empirical analyses are based on rich panel data that come from the first systematic, regulatory collection on large hedge funds, the Securities and Exchange Commission's (SEC) Form PF. Mandated under the Dodd-Frank Wall Street Reform and Consumer Protection Act, our data contain proprietary information on hedge fund assets, including long and short positions in broad asset classes as well as estimates of the aggregate liquidity of these positions. Our data also comprise information on hedge funds' liabilities, such as borrowing, collateral, and investor redemption restrictions. Form PF offers an unprecedented view of the activities of large hedge funds, which often do not report to any of the public databases.[12]

We find that portfolio illiquidity and investor share illiquidity are strongly positively correlated. A cross-sectional regression of average portfolio illiquidity on average investor share illiquidity produces an R-squared greater than 40% when no additional covariates are included. The estimated coefficient is roughly 0.50, which suggests on average share illiquidity increases more quickly than asset illiquidity, and has a t-statistic that exceeds 40. This result offers evidence that at least some of the costs associated with illiquid assets ultimately resides with the hedge fund investor rather than the fund manager.

Next, we estimate the importance of illiquidity for explaining hedge funds' risk-adjusted returns. It is in this context that our data are particularly valuable because our measures of illiquidity span much longer horizons than can be inferred from trading data. We find that average alphas estimated from traditional factor models and based on net returns are 38% smaller after controlling for portfolio illiquidity. Interestingly, investor share illiquidity has considerably more explanatory power for estimated alphas. Investor share illiquidity reduces average estimated alphas by over 64%, and reduces the mean absolute pricing error by 32%. We find similar results using gross returns.

[12]See Barth and Wermers (2018) for a detailed analysis of the differences between the Form PF and Lipper Tass data sets.

Our estimation approach also produces an economically meaningful illiquidity return premium, interpreted as the increase in return due to an additional log-day needed for investors to redeem their shares or for managers to liquidate assets without price impact. Our estimates imply that an additional log-day of investor share illiquidity is associated with an additional excess return of 54 basis points per year XXXX net or gross returns! XXXXX. Alternatively, the *ceteris paribus* excess return earned by the 95^{th} illiquidity-percentile fund is 0.23% higher per month than a fund at the 5^{th} percentile of illiquidity. This compounds to roughly an additional 2.8% rate of return per year. The return premiums estimated from gross and net returns, separately, also give the average fraction of the illiquidity premium earned by investors. Based on our estimates of the gross- and net-return illiquidity premiums, we find that roughly 75% of the illiquidity premium earned by hedge fund assets is passed through to fund investors. To our knowledge, this result has not been established elsewhere in the literature.

Based on the differences in risk-adjusted returns between the most and least liquid funds, we next examine whether our measures of illiquidity — which are properly interpreted as fund characteristics — may also represent an asset pricing risk factor. We construct an illiquidity factor by sorting funds into five groups based on our measures of illiquidity, and calculating the return difference in each period between the most and least liquid portfolios. While this illiquidity risk factor has a 31% correlation with the Pastor and Stambaugh (2003) traded liquidity factor, suggesting it has some information content, we find no evidence that exposure to this factor has explanatory power for the cross-section of fund returns. The estimated price of risk from the illiquidity factor is economically small and statistically insignificant. This indicates that while the illiquidity premium may be risk compensation for fund managers who cannot effectively diversify price impact risks, it is not compensation for exposure to a systematic risk factor.

While the illiquidity risk factor we construct provides no predictive power for the cross-section of hedge fund returns, the risk factor's constituent portfolios do highlight important properties of the returns to liquid and illiquid funds. The returns to the most illiquid funds are both less volatile and more persistent than the returns to liquid funds. One concern is that the smoothness of returns in illiquid portfolios could mechanically reduce estimated betas and therefore increase estimated alphas (Asness, Krail, and Liew (2001), Khandani and Lo (2011)). However, we show that return smoothing as proxied by the first-order autocorrelation coefficient is unlikely to explain our findings. First, the illiquidity premium shows up even in raw, unadjusted returns. Additionally, the largest differences in alphas are found between groups of funds that differ little in their estimated return smoothing, whereas the smallest differences in alphas are found between funds that differ the most in estimated smoothing. Our conclusion is that while return smoothing is an important consideration, it is unlikely to be a primary determinant of our results.

Finally, we explore the relationship between illiquidity, skill, and fees. Consistent with managers retaining some of the risks associated with illiquid assets, we show that realized fees are positively associated with portfolio illiquidity. These results remain after controlling for alpha as an estimate of manager skill, and the proportion of months the fund earns positive gross returns, intended capture a mechanical relationship that results from the manager being more likely to earn the incentive fee and exceed the high-water mark. Further, consistent with investor share restrictions representing a better proxy of the total illiquidity risk, we find that fees are positively associated with investor share illiquidity but not portfolio illiquidity once investor share illiquidity is included in the regression. Higher fees are also associated with higher leverage, suggesting part of the illiquidity risk faced by managers comes from funding liquidity risk (i.e. margin risk on leveraged positions). Further, the relationship between realized fees and illiquidity is entirely driven by the incentive fee, which is consistent with managers capturing more of the illiquidity premium as compensation for undiversifiable risk.

We also find that while alphas estimated from gross returns are strongly associated with both the management and incentive fees, alphas estimated from net returns are uncorrelated with fees. This provides empirical evidence that competition in the hedge fund industry drives net-of-fee alphas to zero, consistent with the theory of Berk and Green (2004). Such an exercise is only possible with measures of both gross- and net-of-fee returns of hedge funds. To our knowledge, this finding is also new to the literature.

Our paper is most closely related to Aragon (2007) and Khandani and Lo (2011). Aragon (2007) shows that investor share restrictions can explain the entirety of average hedge fund alpha using a sample of hedge funds from the Lipper TASS database. That paper differs from ours in a few important ways. First, the TASS database primarily comprises smaller funds; our data on large funds is much more representative based on total assets under management. Second, Aragon (2007) uses investor share restrictions as a proxy for portfolio illiquidity. However, we show that even with estimates of the average liquidation horizon of hedge fund assets, investor share restrictions do significantly better at explaining estimated alphas. Such a result is only possible when both investor share illiquidity and the illiquidity of the underlying assets are observable. Finally, using data on gross and net returns, we are able to determine how much of the illiquidity premium is ultimately captured by investors. Khandani and Lo (2011) estimate the illiquidity premium in hedge funds using the first-order autocorrelation coefficient as a measure of return smoothing and a proxy for illiquidity, a method based on Getmansky, Lo, and Makarov (2004).[13] They estimate a premium of 3.96%

[13]The Getmansky, Lo, and Makarov (2004) method of estimating illiquidity from smoothness in returns is by far the most common approach in the hedge fund literature: Bollen and Pool (2008), Ding, Shawky, and Tian (2009), and Aiken, Clifford, and Ellis (2015) are just a few examples of many that use these measures of return serial correlation to estimate portfolio illiquidity levels in hedge funds.

across all funds, and much larger premiums within strategy. We find, however, that return smoothness is only partially correlated with illiquidity in our sample, and that return premiums estimated from autocorrelations in our Fama-MacBeth procedure yield statistically insignificant return premiums with the wrong sign.

Our results are related to other studies of hedge fund liquidity, but differ in important ways. Sadka (2010) shows that funds with returns that more tightly covary with unexpected changes in aggregate liquidity have better performance on average, but underperform during times of low liquidity. Teo (2011) shows that funds with more generous redemption policies and higher exposures to market-wide liquidity risk are more likely to sell assets at discounted prices to meet investor redemptions. Cao, Chen, Liang, and Lo (2013) estimate whether funds successfully time market liquidity by adjusting their market betas. Our paper differs from these in that we examine the return premium associated with illiquidity levels rather than the covariance with measures of aggregate market liquidity.

The paper is organized as follows: Section 2 provides an economic argument for the reduced-form empirical relationship between hedge fund returns, portfolio illiquidity, investor share illiquidity, and fees. Section 3 describes our data and provides fund-level summary statistics. Section 4 documents hedge fund illiquidity summary statistics and establishes the empirical relationships between portfolio illiquidity and investor share illiquidity. Section 5 evaluates the importance of illiquidity in explaining hedge fund alphas and provides estimates of the illiquidity return premium. Section 6 analyzes the relationship between manager skill, illiquidity, and fees. Section 7 concludes.

2 An Economic Argument for Illiquidity, Share Restrictions, and Fees

This paper uses new regulatory data on large hedge funds to estimate the return premium earned by illiquid assets and the fraction of the premium captured by hedge fund investors. Specifically, we empirically estimate the size of the illiquidity premium, as well as the relationship between investor share restrictions and portfolio illiquidity, and between portfolio illiquidity and fees.

In this section, we provide context for these empirical relationships by outlining an economic mechanism that links asset illiquidity, expected returns, investor share restrictions, and fees. While our empirical results only establish reduced-form relationships in the data, and are therefore not tied to a specific model, the economic argument provided here motivates our findings and endows them with an economic interpretation. A formal theoretical model is outside the scope of this paper, but is an interesting area for future research.

2.1 The Risks of Illiquidity in Hedge Fund Portfolios

Modern asset pricing theory establishes that expected returns on risky securities are determined by their sensitivity to undiversifiable risk factors. Valid risk factors are constructed as returns to traded portfolios that represent a systematic risk.[14] However, recent research in the limits to arbitrage suggests that some risks that are in principle diversifiable cannot be diversified in practice, and therefore require compensation in equilibrium. Examples of such risks include prepayment risk on mortgage-backed securities (Gabaix, Krishnamurthy, and Vigneron (2007)) and shocks to intermediary capital (He and Krishnamurthy (2012), Siriwardane (2018)). We argue that in delegated asset management, the risk associated with price impact from forced sales of specialized or thinly traded assets is also a risk that requires compensation in equilibrium.

In markets with specialized and difficult to price assets, the costs of information acquisition and the necessity of asset-specific expertise give rise to natural barriers to entry. This implies the marginal investor is likely to be a specialist rather than a well-diversified trader. This also implies that fewer traders exist in such markets, and if assets need to be sold quickly there are fewer buyers willing to buy such assets at fair-value prices. Selling specialized assets may be particularly difficult if the liquidity shocks faced by one specialist are likely to be experienced by other specialists in the same market. Specialists who are forced to sell assets quickly may therefore be required to offer price concessions to less-sophisticated diversified traders or other similarly constrained specialists. This is the foundation of the fire sale risk mechanism outlined in Shleifer and Vishny (1992).

For the specialist to be sufficiently compensated for his asset-specific expertise, he will likely need to raise outside capital. Fees on outside investors' capital effectively allow the specialist to leverage the returns on his personal wealth and reduce the required rate of return in specialized markets. Outside investors, who have less trading skill or have made investments in generalized rather than asset-specific human capital, can only invest in specialized markets through specialists. Further, because public funds such as U.S. mutual funds are restricted in the amount of capital that can be allocated to illiquid assets, traders in illiquid markets are likely to be private fund managers, such as endowments or hedge funds. In this paper, we focus on hedge fund managers as specialists who trade in illiquid markets, and are therefore assumed to be the marginal investor in such markets.

An important feature of hedge funds is that hedge fund managers invest substantial portions of their personal wealth in the fund. This serves to align the incentives of fund managers and fund investors, who

[14]The imposition that factors are constructed based on *traded* securities ensures that the arbitrage mechanism in financial markets will eliminate all return differentials that emanate from non-risk sources.

are unable to fully monitor the trading and pricing activities of the fund due to the opacity and complexity of hedge fund strategies. But fund managers' significant investment of personal capital also implies that managers are unable to fully diversify the risk of price impact due to immediate liquidity needs. If investor redemptions or margin calls force the fund manager to sell assets at a discount from the price marked in the fund's net asset value, it is the manager and remaining investors who bear this cost as the residual claimants on fund assets. In fact, the manager is the ultimate residual claimant; if all investors redeem, the manager will remain as the lone fund investor. Because a significant portion of the manager's personal wealth is invested in the fund, and therefore in a single asset class or strategy, and because the manager is unable to redeem his personal capital without closing the fund entirely, the idiosyncratic risks that arise in these assets constitute an undiversifiable risk for the fund manager.[15]

2.2 Empirical Predictions

For managed funds, immediate liquidity needs originate from two primary sources: funding liquidity shocks associated with margin calls or expiring debt on leveraged investments, or from investor redemptions. If hedge fund managers are unable to diversify the illiquidity risk associated with specialized, thinly traded assets, they are left with two options: require explicit compensation for this risk in the form of higher fees, or pass some of the potential costs of price impact to fund investors in the form of share restrictions (or both).

To formalize ideas, imagine that the cross-section of fund returns are described by a standard asset pricing model of the form estimated in Section 5:

$$E[R_j^{gross}] = \alpha_j + \lambda L_j + \gamma' \beta_j \tag{1}$$

The model assumes that the expected return to the portfolio of fund j is determined by sensitivities to standard risk factors (β_j) and the prices of risk for each factor (γ), the illiquidity of fund assets L_j and the illiquidity premium (λ), and an unexplained component α_j. Note that L_j is a measure of illiquidity, so a higher value indicates the asset is more illiquid. In Section 3, we precisely define our measures of asset and investor share illiquidity, which are the weighted-average number of days needed to sell assets without price impact, and the weighted-average number of days until investors' capital is returned once a redemption is requested.

Further, suppose that the fee the hedge fund manager charges is some fraction of net assets under management, $\pi_{j,0}$, and some proportion of the return on fund assets, $\pi_{j,1}$. Both $\pi_{j,0}$ and $\pi_{j,1}$ are percentages.

[15] Some funds operate as "multi-strategy" funds, investing in a range of strategies and asset classes. Nonetheless, such funds comprise only a portion of total hedge fund assets, and are not fully diversified across the entirety of possible asset classes.

Then, we can write the net return on assets earned by investors as:

$$E[R_j^{net}] = (1-\pi_{j,1})\big(\alpha_j + \lambda L_j + \gamma' \beta_j\big) - \pi_{j,0} \tag{2}$$

First, consider the case in which the manager passes *all* of the illiquidity costs to fund investors through the imposition of share restrictions. Share restrictions often take the form of an initial lockup period, a redemption notice period that establishes a number of days between the initial investor redemption request and when the fund actually meets those redemptions, and a redemption frequency that establishes how often funds will pay redemptions (for example, quarterly). This implicitly assumes that there is no financing risk associated with margin on leveraged positions, and that the only funding liquidity risk managers face is from investor redemptions.

In this case, competition will will force the entirety of the illiquidity premium (λ) to be passed to fund investors, and the incentive fee $\pi_{j,1}$ will be a function of L_j. For simplicity, assume illiquidity is the only source of fee heterogeneity: $\pi_{1,j} = \pi_1(L_j)$ and $\pi_{0,j} = \pi_0$. We then have:

$$\frac{\partial E[R_j^{net}]}{\partial L_j} = \frac{\partial E[R_j^{gross}]}{\partial L_j} = \lambda \tag{3}$$

$$\Longrightarrow \frac{\partial \pi_1(L_j)}{\partial L_j} = -\frac{\lambda \pi_1(L_j)}{\lambda L_j + \gamma' \beta_j + \varepsilon_j} \tag{4}$$

If $E[R_j^{gross}] > 0$, $\lambda > 0$, and $\pi_1 > 0$, meaning that expected gross returns, the illiquidity premium, and incentive fees are all positive (as one would expect), then for the entirety of the illiquidity premium to be passed to fund investors it must be that incentive fees and illiquidity are negatively correlated. That is, $\frac{\partial \pi_1(L_j)}{L_j} < 0$. Further, as the illiquidity of the assets increases so too do the share restrictions; this is a necessary condition for managers to continue to ensure that assets will never need to be sold at a discount. This defines an empirically testable hypothesis:

> *If fund managers pass the entirety of illiquidity risk to fund investors: (1) asset illiquidity and investor share restrictions will be positively correlated; (2) the illiquidity premium (λ) will be identical for gross and net returns; and (3) the incentive fee will be negatively correlated with asset illiquidity.*

Alternatively, consider the case in which the manager imposes no share restrictions, or in which share restrictions originate from economic forces that are independent of illiquidity. In this case, the manager is exposed to the full amount of undiversifiable illiquidity risk that results from investors' redemption options, and will therefore require additional compensation as a result. Because the investors' shares are unrestricted (or are restricted in a manner that is uncorrelated with illiquidity), the manager will capture the entirety of the illiquidity premium. This implies:

$$\frac{\partial E[R_j^{net}]}{\partial L_j} = 0 \tag{5}$$

$$\implies \frac{\partial \pi_1(L_j)}{\partial L_j} = \frac{\lambda(1-\pi_1(L_j))}{\alpha_j + \lambda L_j + \gamma' \beta_j} \tag{6}$$

Under the conditions $E[R_j^{gross}] > 0$, $\lambda > 0$, and $\pi_1 > 0$, it follows that $\frac{\partial \pi_1(L_j)}{\partial L_j} > 0$. That is, if the manager captures the entirety of the illiquidity premium, fees and illiquidity must be positively correlated all else equal. This offers an alternative empirical hypotheses:

> *If fund managers pass none of the illiquidity risk associated with investor redemptions to fund investors: (1) asset illiquidity and investor share restrictions will be uncorrelated; (2) the illiquidity premium (λ) will be evident only in gross returns; and (3) the incentive fee will be positively correlated with asset illiquidity.*

Finally, if some but not all of the illiquidity risk associated with investor redemptions is passed to investors, or if illiquidity risk in part results from risks associated with margin on leveraged positions (which cannot be passed to investors), then managers and investors will share the costs of illiquid assets as well as the illiquidity premium. This yields a set of empirical predictions that fall between the first two hypotheses:

> *If fund managers pass some (but not all) of the illiquidity risk associated with investor redemptions to fund investors, or if illiquidity risk results in part from the risk associated with margin on leveraged positions: (1) asset illiquidity and investor share restrictions will be positively correlated; (2) the illiquidity premium (λ) will be partially but not fully reflected in net returns, and the ratio of the illiquidity premium estimated from gross and net returns will determine the fraction of the illiquidity premium captured by fund investors, and (3) the incentive fee will be more positively correlated with asset illiquidity than if the manager passed all of the illiquidity risk to fund investors.*

We note that it is possible for fees and illiquidity to be negatively correlated even if the manager passes some of the risk to fund investors (the correlation just can't be *too* negative). However, a sufficient condition to rule out the first hypothesis that the manager retains none of the risks associated with illiquid assets is that fees and illiquidity are positively correlated.

One issue yet to be addressed is whether investors can fully diversify the costs associated with illiquid assets. If they can, all of the illiquidity risk to managers should be passed to investors, and illiquid assets held by hedge funds should not earn a return premium. To investigate this, imagine a hedge fund investor that is diversified across many different hedge funds that invest in different illiquid asset classes. Further, imagine the price impact associated with these assets is uncorrelated, so that full illiquidity diversification may be possible. Yet, if managers pass the risks arising from investor redemptions to fund investors by restricting their shares, the investor is unable to redeem his capital quickly from any of them. That is, the investor's capital is fully locked up for a given period of time even though he is diversified across many assets. This is because the investor can only invest in illiquid assets through the hedge fund specialist as an intermediary. Because the investor would prefer to be able to redeem capital on demand, he will only invest in such funds if the expected returns from investing in an illiquid fund exceed those from a liquid fund. Because in this case the hedge fund investor becomes the marginal investor in illiquid asset markets rather than the manager, the underlying illiquid assets will still command a return premium. This argument suggest that no matter who ultimately bears the costs of illiquid assets, the fund manager or fund investor, illiquid assets will earn a premium in equilibrium.

Finally, we speculate that the extent to which illiquidity risk is passed to investors versus held by managers will ultimately be determined by the relative risk aversions of the managers and investors. If investors bear more of the illiquidity costs it must be that they are better able to diversify them across various asset classes. This is likely, as fund managers have a significant portion of their personal wealth in the fund, whereas fund investors are able to invest across a wider variety of funds and asset classes. Unfortunately, this is not an empirically testable statement given our data, and we leave such analyses for future research.

2.3 Interpretation of Illiquidity "Risk"

The term "risk" is often reserved for exposure to systematic risk factors. In this paper, illiquidity risk refers to the undiversifiable risk faced by fund managers with a large personal stake in the fund, and therefore in a specific asset class or strategy. The argument that fund managers will require compensation for this risk is a unifying framework that provides coherence to the novel empirical results we document in Sections 4.3-6. We are therefore tempted to conclude that the economic argument above is the appropriate paradigm for understanding illiquidity premium, share restrictions, and fees in hedge funds.

However, the deterioration of returns associated with price impact is usually described as "transactions costs". While it is unclear why assets with higher transactions costs would be associated with higher fees (an empirical finding discussed in Section 6), our empirical results are nonetheless unchanged by the interpretation that the illiquidity premium represents transactions costs rather than a compensated risk in delegated asset management. Further, while illiquid assets may pose risks to undiversified fund managers, the negative impact of share restrictions on investors *should* be considered a transactions cost rather than a risk. This is because fund investors are the principal in delegated asset management. Thus, throughout the remainder of the paper, when discussing the negative effects of share restrictions on investors we will refer them as illiquidity costs, and reserve the term illiquidity risk only for discussions of fund managers.

Because our measures of illiquidity are fund characteristics rather than risk factors, we also refrain from using the term "price of risk" in our asset pricing specifications, choosing instead the term "premium" to describe the higher risk-adjusted returns associated with illiquid assets. While higher expected returns may originate from required compensation from asset managers, none of our empirical tests can definitively rule out alternative hypotheses for the ultimate source of illiquidity premia.

3 Data

We use fund-level data from the Security and Exchange Commission's Form PF (PF for "Private Fund"), which is the first systematic regulatory collection of hedge fund data in the United States. Form PF was adopted in a joint rule-making by the SEC and Commodity Futures Trading Commission in 2011 to fulfill a mandate of the Dodd-Frank Wall Street Reform and Consumer Protection Act of 2010. The form is filed by investment advisers registered with the SEC who manage at least $150 million in private funds, such as hedge funds and private equity funds.[16] Small private fund advisers file annually and report items such as gross and net asset values, returns, total borrowings, strategies, investor composition, and their largest counterparties. *Large Hedge Fund Advisers* — those with at least $1.5 billion in assets managed in hedge funds — are required to report this information at a quarterly frequency as well as more detailed information regarding portfolio, investor, and financing illiquidity, asset class exposures, collateral posted and risk metrics, and more, for each of their *Qualifying Hedge Funds*.[17] A Qualifying Hedge Fund has a net asset value (NAV) of at least $500 million as of the last day in any month in the fiscal quarter immediately preceding the adviser's most recently completed fiscal quarter.[18] In this paper, we focus exclusively on Qualifying Hedge

[16]For details on who must register as an investment adviser with the SEC, see https://www.sec.gov/about/offices/oia/oia_investman/rplaze-042012.pdf.

[17]This is reported in Section 2b of Form PF.

[18]The thresholds for filing Form PF and for the Large Hedge Fund Adviser classification are on a gross basis, but the threshold for Qualifying Hedge Fund status is on a net basis. Moreover, when determining whether a reporting threshold is met, advisers must aggregate the asset values of the funds themselves, associated parallel funds, dependent parallel managed accounts, and master-feeder funds. Advisers must also include these items for related persons that are not separately operated. Finally, while reporting thresholds are determined on an aggregated basis, advisers are permitted to report fund-level data on either an aggregated or disaggregated basis. Thus, some qualifying hedge funds in our sample have NAV less than the associated threshold of $500 million. See Flood, Monin, and Bandyopadhyay (2015) and Flood and Monin (2016) for more information on the structure and history of Form PF.

Funds.[19]

We apply several filters to the raw data to form our main analytic sample. Due to data quality concerns at the beginning of the Form PF data collection, we restrict our data to filings from 2013 and after. We require that funds have at least one non-missing observation for each of our primary variables of interest and our standard set of controls, which include data on portfolio illiquidity, investor illiquidity, financing illiquidity, borrowing, net asset value (NAV), gross asset value (GAV), and asset class exposures. We further require gross assets to be greater than or equal to net assets. Funds are also asked to report the fraction of assets priced by three criteria (referred to as Level 1, 2 and 3 assets), and we require the values of each of these categories be non-negative.

Whereas the majority of our data are generally reported on a quarterly basis, funds report returns on a monthly basis. Rather than rates of return, some funds appear to report an internal rate of return (IRR) since inception. In order to exclude funds that may be reporting IRRs instead of rates of return, we exclude funds with Sharpe ratios greater than 2. Our resulting base sample is an unbalanced panel of 2,579 funds for time periods between January 2013 and September 2017. In total, we have 92,148 fund-month observations and 34,198 fund-quarter observations (for quarterly filers).

The top panel of Table 1 reports cross-sectional summary statistics for our main sample. Funds have an average net asset value of $1.13 billion and a median of $0.51 billion, which suggests that fund size is positively skewed. Skewness is also present in the distribution of gross asset value, which averages $2.04 billion but has a median of $0.68 billion. Balance sheet leverage, defined as the ratio of gross to net assets, is on average 3.43 but has a median and 75th percentile of 1.16 and 1.67, respectively, suggesting a high degree of skewness. Over the 19 quarters in our sample, we have on average 13.26 quarterly observations per fund and 35.73 monthly returns observations.

It is useful to distinguish the data in this paper from data reported to public hedge fund data services, such as Lipper TASS, HFR, or Morningstar. Public databases are known to largely contain smaller funds, which use these services to market and raise investor awareness. For example, as of 2016 the assets under management for the median fund reporting to TASS is around $30 million. The 90th percentile fund in TASS has just under one billion dollars in assets under management. Thus, this paper represents the first empirical analysis of the illiquidity characteristics of large hedge funds, which comprise the vast majority of total hedge fund assets.

The bottom panel of Table 1 contains counts of funds by strategy. Our strategy categorizations are

[19]Form PF data are confidential. The Office of Financial Research has access to the data through an agreement with the SEC. The form itself is publicly available and can be downloaded here: https://www.sec.gov/rules/final/2011/ia-3308-formpf.pdf.

determined by funds' responses to 22 pre-selected strategy categories, which fall within eight broader classifications: Credit, Equity, Event Driven, Macro, Relative Value, Managed Futures, Invests in Other Funds, and Other.[20] A hedge fund is categorized as pursuing a given broad strategy if 75% or more of its normalized assets are allocated to that strategy. If there is no broad strategy category that contains 75% or more of the fund's normalized assets, we classify the fund as multi-strategy. Table 1 shows Equity is the most common strategy, with 807 funds, while Managed Futures and Invests in Other Funds are the least common strategies. We note the large contingent of funds that report their strategy as "Other", and do not fall into any of the primary strategy classifications. Finally, in order to avoid double counting of hedge fund assets, Form PF asks funds not to report information for investments in other funds; for this reason, we exclude funds categorized as following an "invests in other funds" strategy.

Finally, because Form PF data is confidential and highly sensitive, we aggregate observations to mask information about any individual filer. Therefore, none of our results are provided at the individual fund level. When a sufficient level of aggregation is not possible, we replace reported values with a star.

4 The Illiquidity of Assets and Investor Shares

Section 2 offers an economic rationale for the empirical relationships between asset illiquidity, investor share restrictions, risk-adjusted returns, and fees. In this section, we develop our measures of portfolio illiquidity and investor share illiquidity based on estimated time horizons reported by funds themselves. We then describe the properties of these illiquidity measures, and estimate the extent to which these measures are associated.

4.1 Portfolio Illiquidity

Our measure of portfolio illiquidity is derived from hedge funds' responses to estimated liquidation horizons. Specifically, the form asks funds to:

> "Provide the following information regarding the liquidity of the reporting fund's portfolio. Specify the percentage by value of the reporting fund's positions that may be liquidated within each of the periods specified below. Each investment should be assigned to only one period and such assignment should be based on the shortest period during which you believe that such position could reasonably be liquidated at or near its carrying value. Use good faith estimates for liquidity based on market conditions over the reporting period and assuming no fire-sale discounting. In the event that individual positions are important contingent parts of the same trade, group all those positions under the liquidity period of the least liquid part (so, for example, in a convertible bond arbitrage trade, the liquidity of the short should be the same as the convertible bond. Exclude cash and cash equivalents."

[20] This is question 20 on Form PF.

The periods specified are 0-1 days, 2-7 days, 8-30 days, 31-90 days, 91-180 days, 181-365 days, and more than 365 days. For each fund i and quarter t, we construct our primary measure of portfolio illiquidity, $L_{p,i,t}$:

$$L_{p,i,t} = \sum_{j=1}^{7} \pi_{j,i,t} \overline{H}_j, \tag{7}$$

where $\pi_{j,i,t}$ is the fraction of hedge fund $i's$ portfolio that can be liquidated in period j and quarter t, and $\overline{H}_j$ is the midpoint of period j. For example, $\overline{H}_1 = 0.5$ days. For the most illiquid bin (greater than 365 days), we assume that $\overline{H}_7 = 366$. $L_{p,i,t}$ is simply the weighted-average number of days needed to liquidate fund assets without price impact. It is not the number of days needed to liquidate the entire portfolio at fair value prices, as that is determined by the number of days needed to sell the most illiquid asset. We also emphasize that $L_{p,i,t}$ is an *expectation* of the true weighted-average liquidation horizon, which is a random variable.

We also compute the within-fund average of $L_{p,i,t}$, denoted by $\overline{L}_{p,i}$. That is,

$$\overline{L}_{p,i} = \frac{1}{T_i} \sum_{t=1}^{T_i} L_{p,i,t},$$

where T_i is the number of quarters for which we observe fund i. For ease of notation, we omit the i and t subscripts of $L_{p,i,t}$ and $\overline{L}_{p,i}$ when the context is clear.

Figure 1 shows the distribution of $\overline{L}_p$ for our sample of 2,579 hedge funds. The majority of assets held by funds can be liquidated in less than seven days. This is consistent with equities, macro, and managed futures strategies that trade in equities, derivatives, and other liquid assets. However, the distribution of portfolio illiquidity has a long right-tail; 907 funds report an average value of portfolio illiquidity that is greater than 30 days, and 513 report a value greater than 100 days or more. Table 2 shows the mean, median, and standard deviation of $\overline{L}_p$, both in aggregate and separately for each broad hedge fund strategy. Unsurprisingly, Equity and Macro strategies have the lowest $\overline{L}_p$, with medians of 6.3 and 2.9 days, respectively. The most illiquid strategies are Event Driven and Credit, with median values of 41.3 and 28.1, respectively.

There are a few important caveats associated with the measure L_p. First, L_p measures the expected time needed to liquidate the hedge fund's *entire* position in that asset, rather than the time needed to liquidate each individual share of a given security. For example, if the fund holds 100,000 shares of security A, and can liquidate 50,000 shares in one day without price impact, but would need seven days to liquidate the full 100,000-share position without price impact, the fund will report a seven day liquidation horizon for all 100,000 shares. Secondly, funds report expected liquidation periods assuming normal market conditions, but such liquidation periods may be considerably longer during times of low liquidity or market stress. Finally, while the liquidation periods offered are rather tight for high levels of liquidity (0-1 days, 2-7 days), these periods widen considerably for less liquid assets. An asset that can be sold at fair value in 31 days is substantially more liquid than one that would need 89 days, but both assets would end up reported in the same liquidation period: 31-90 days. This will introduce some noise in our estimates of the most deeply illiquid assets.

4.2 Investor Share Illiquidity

Similarly to the portfolio illiquidity measure L_p, our data allow us to estimate an analogous measure of investor share illiquidity. Specifically, the form asks funds to:

> "Divide the reporting fund's net asset value among the periods specified below depending on the shortest period within which investors are entitled, under the fund documents, to withdraw invested funds or receive redemption payments, as applicable. Assume that you would impose gates where applicable but that you would not completely suspend withdrawals/redemptions and that there are no redemption fees. Please base on the notice period before the valuation date rather than the date proceeds would be paid to investors".

The time periods provided are the same as those supplied in the portfolio illiquidity question. Just as with portfolio illiquidity, we construct a measure of investor share illiquidity. For each fund i and quarter t, define $L_{s,i,t}$ by

$$L_{s,i,t} = \sum_{j=1}^{7} \psi_{j,i,t} \overline{W}_j, \tag{8}$$

where $\psi_{j,i,t}$ is the fraction of hedge fund i's shares that can be redeemed in period j and quarter t, and $\overline{W}_j$ is the midpoint of period j. As with L_p, for the most illiquid bin (greater than 365 days) it is assumed that $\overline{W}_7 = 366$. We also compute the within-fund average of $L_{s,i,t}$, denoted by $\overline{L}_{s,i}$. That is,

$$\overline{L}_{s,i} = \frac{1}{T_i} \sum_{t=1}^{T_i} L_{s,i,t},$$

where T_i is the number of quarters for which we observe fund i. As before, we omit the i and t subscripts of $L_{s,i,t}$ and $\overline{L}_{s,i}$ when the context is clear.

The left-hand panel of Figure 2 shows that investor share illiquidity is of considerably longer duration than reported portfolio illiquidity. While a surprising number of funds offer on-demand redemptions (roughly 12% of funds in our sample offer seven day or less redemptions), most funds impose non-trivial share restrictions on investors. Table 2 shows that the median value of $\overline{L}_{s,i}$ is 135.5 days. The most illiquid shares are found in Event Driven funds, which have a median investor share illiquidity of 283.6 days.

4.3 The Relationship Between Investor Share Restrictions and Portfolio Illiquidity

The right-panel of Figure 2 shows a bin-scatter plot of average investor share illiquidity versus average portfolio illiquidity. The bottom panel of Figure 2 shows a bin-scatter plot of log average portfolio illiquidity, $\log(\overline{L}_p)$, and average investor share illiquidity, $\overline{L}_s$. The figure shows that portfolio and investor share illiquidity are strongly positively related. For the vast majority of log portfolio illiquidity values, the relationship appears nearly linear. The figure is highly similar if medians are used rather than means. This suggests that the level-level relationship between investor share illiquidity and portfolio illiquidity is highly concave; investor share restrictions increase much more quickly than portfolio illiquidity as the portfolio shifts from highly liquid to somewhat less liquid assets, but increases more slowly once the portfolio primarily comprises illiquid assets. The relatively slow increase in investor share illiquidity at increasingly illiquid portfolios may in part be because for these portfolios, investor share restrictions are already highly restricted, and our measure of investor share illiquidity has a maximum value of $L_s = 366$.

In Table 3, we explore the relationship between investor share illiquidity and portfolio illiquidity more formally. Column (1) shows the results from a cross-sectional regression of the average portfolio illiquidity on average share illiquidity, with no additional controls. A coefficient of 0.499 suggests that each additional day of portfolio illiquidity is associated with two more days of investor share illiquidity. The coefficient is highly statistically significant. Further, the R-squared in this regression is over 40%, which suggests much of the variation in portfolio illiquidity can be explained by variation in investor share illiquidity. In column (2) we add additional controls for size, balance sheet leverage measured as the ratio of gross to net assets, and broad strategy, and find the coefficient changes only slightly. In column (3), we use the log of average portfolio illiquidity based on the log-linear relationship shown in Figure 2, and find that the model fit is somewhat improved.

In columns (4) and (5) of Table 3, we include time and fund fixed effects (separately). The results from the specification with time fixed effects show that the coefficient is highly similar to the coefficient in the cross-sectional regressions. In the specification with fund fixed effects, however, the coefficient is dramatically reduced, and is no longer significant at the 1% level. This indicates that the covariation between portfolio and investor share illiquidity operates through heterogeneity in fund characteristics, rather than through time variation *within* a fund. This is consistent with a model in which funds set share restrictions based on their expected portfolio holdings, but do not alter redemption restrictions based on time-variation in the liquidity of their assets. Share restrictions are therefore best thought of a fund fixed effect, rather than a time-varying risk-management tool.[21] These results are consistent with those found in Aragon, Ergun, Getmansky, and Girardi (2017b), who also investigate the relationship between portfolio, investor, and financing liquidity using Form PF data.

In the context of the economic argument outlined in Section 2, the tight relationship between investor share illiquidity and portfolio illiquidity offers evidence that at least some of the illiquidity risk associated with the potential price impact from forced asset sales is passed from the manager to fund investors. Such an outcome would make sense if hedge fund investors could more easily diversify their investments across various asset classes. Given the large investments of personal wealth in hedge funds by hedge fund managers, we view this as a reasonable possibility.

5 The Illiquidity Premium

Section 2 offered an economic rationale for an empirical relationship between illiquidity, expected returns, investor share restrictions, and fees. Consistent with fund managers passing at least some of the risk associated with illiquid assets to investors, section 4.3 showed that investor share restrictions are strongly correlated with the illiquidity of the underlying assets. Section 2 also argued that regardless of who bears the cost of illiquid assets, such assets should command a return premium in equilibrium. In this section, we estimate the size of the illiquidity premium and its importance for explaining the cross-section of hedge fund returns.

Estimating the size of the illiquidity premium is difficult. In virtually all previous studies of illiquidity, measures are based on data from trades. Such measures include price impact per volume (Amihud and Mendelson (1986), Amihud (2002), Sadka (2006)) or price reversals (Pastor and Stambaugh (2003)). But assets that trade enough for such measures to be precisely estimated are, by construction, the most liquid securities. Many assets trade relatively infrequently, or in some cases not at all. The premium earned by such deeply illiquid assets cannot be measured by assets that trade often enough for previous liquidity measures to be well-defined.

The data described in section 3 offer a way forward. Instead of estimating the illiquidity associated with individual assets and calculating their expected returns, we can instead examine the expected returns of portfolios of illiquid assets over time. Because the illiquidity measures we examine in this paper relate to the time needed to sell assets without price impact, or for investors to redeem shares, we can use these to estimate the premium earned over a wide range of asset illiquidity, not just the premium earned by the most-often traded illiquid assets.

[21] For a detailed analysis of hedge fund liquidity risk management practices using Form PF data, see Aragon, Ergun, Getmansky, and Girardi (2017a).

In our empirical analysis of the illiquidity premium, we employ a two-stage model that produces a panel of pricing errors, the average of which (within fund) we interpret as manager skill and which we use later in our analysis of hedge fund fees. We then examine the extent to which illiquidity can explain these pricing errors, which in the absence of our measures of illiquidity would be incorrectly interpreted as manager skill. Our approach also produces an estimate of the illiquidity premium that includes deeply illiquid assets, which has been impossible to estimate previously.

5.1 The Two-Stage Pricing Model

Our estimates of the cross-section of hedge fund returns follow the standard two-stage regression approach of Fama and MacBeth (1973). In the first stage, we estimate betas on traded factors from time-series regressions run separately fund-by-fund:

$$R_{i,t} = \beta_i' f_t + \varepsilon_{i,t}, \tag{9}$$

where f_t is the set of traded factors and includes a vector of ones, so that the vector β_i includes an intercept. We estimate β_i in (9) from three factor models. The first is the Fama-French-Carhart four-factor model (FF), which includes the three traditional Fama-French factors and a momentum factor (Fama and French (1992), Carhart (1997)). The second is the Fung-Hsieh seven factor model (FH), designed explicitly to represent common risk factors in hedge fund returns (Fung and Hsieh (2001), Fung and Hsieh (2004))[22]. The FH model includes risk-factors for stock market exposures, interest rates, credit spreads, and trend-following factors for currency, commodity, and bond markets.[23] Finally, our third model combines the seven FH factors with the three Pastor-Stambaugh liquidity factors (FH+PS), which comprise market-wide measures of liquidity (Pastor and Stambaugh (2003)) and ensure that our measures of illiquidity are orthogonal to factors based on time-varying, systematic market liquidity.

In much of our analysis, the (FH+PS) model will be our primary specification, but results are quantitatively similar for the other models. We denote the factor loadings estimated from these models as: $\widehat{\beta}_i^{FF}$, $\widehat{\beta}_i^{FH}$, and $\widehat{\beta}_i^{FH+PS}$ respectively. The intercepts estimated in (9), β_0, provide an one estimate of hedge fund alpha, and are denoted β_0^{FF}, β_0^{FH}, and β_0^{FH+PS}.

Table 4 shows summary statistics for raw monthly returns, both in aggregate and by strategy, as well as the intercepts estimated from the first-stage regressions in equation (9) for the three different factor models.

[22] The Fung-Hsieh factor data can be accessed at http://faculty.fuqua.duke.edu/~dah7/DataLibrary/TF-FAC.xls

[23] Specifically, the seven Fung-Hsieh factors are: the excess return on the S&P 500; a small-minus-big factor constructed as the difference in return on the Russell 2000 Index and the S&P 500; the change in the Moody's Baa yield minus the change in the 10-year constant maturity Treasury note yield; and the returns on portfolios of lookback straddle options on currencies, commodities, and bonds.

The intercepts in equation (9) represent one measure of hedge fund alphas (Jensen's alpha). Overall, the data show that qualifying hedge funds earn just over 50 basis points per month on average, with Equity being the most lucrative strategy at 68 basis points per month, followed by Multi-strategy and Event Driven, each earning 48 basis points per month on average. Macro had the worst performance at 22 basis points per month. Funds also show substantial Jensen's alphas; the average monthly alpha ranges from 20-35 basis points, depending on the model, and Relative Value and Credit strategies earn roughly 40 basis points each. The lowest alphas come from the Macro and Equity strategies, which nonetheless deliver alphas that exceed ten basis points per month in almost every case. In total, hedge funds in the sample appear to earn substantial risk-adjusted returns.

In the second stage, we run cross-sectional regressions of returns $R_{i,t}$ (net or gross of fees, depending on the specification) on the $\widehat{\beta}_i$ estimated in the first stage and on our measures of illiquidity, separately for each month t. Specifically, we estimate second-stage models using either $\log(L_p)$ or $\log(L_s)$. Because illiquidity data are reported only quarterly, whereas fund returns are reported monthly, we fill illiquidity measures forward to populate data in non-quarter-end months. That is, if illiquidity is reported in month t, we assume $L_{i,t} = L_{i,t+1} = L_{i,t+2}$. This gives us a monthly panel of illiquidity, betas, and fund returns.

5.2 A Baseline Specification without Illiquidity Risk

To establish a baseline, we first estimate the second stage excluding our illiquidity measures, which produces a cross-sectional distribution of hedge fund alphas that ignores illiquidity:

$$R_{i,t} = \phi_t' \widehat{\beta}_i + a_{i,t}, \tag{10}$$

$$E[R_{i,t}] = \widehat{\phi}' \widehat{\beta}_i + \widehat{\alpha}_i, \tag{11}$$

$$\widehat{\phi} = \sum_t \widehat{\phi}_t / T, \;\; \widehat{\alpha}_i = \sum_i \widehat{a}_{i,t} / T, \tag{12}$$

where (10) includes an intercept and $a_{i,t}$ is the pricing error for fund i in month t. The Fama-MacBeth measure of hedge fund alpha is then the time-series average of the cross-sectional pricing errors estimated in each t: $\widehat{\alpha}_i = \sum_t \widehat{a}_{i,t} / T$.

As a baseline specification, the FH+PS model appears to do a good job of explaining the cross-section of hedge fund returns. Figure 3 shows the R-squared for each cross-sectional regression in the second stage of the Fama-MacBeth procedure. Both the mean and median R-squares are around 43%, with a high value above 80%, although the factors appear to be less successful later in the sample. Nonetheless, Figure 3 provides evidence that a standard factor model comprised of the FH+PS factors is an appropriate model for evaluating manager skill in large hedge funds.

Table 5 shows the distribution of the coefficients estimated in equation (10), as well as the t-statistics computed using the sample mean and standard errors across the T individual regressions. The first four factors have positive estimated prices of risk, although none are statistically significant at the 10% level. Two of the three liquidity factors are either statistically significant or nearly significant at the 10% level, but each has a negative estimated price of risk. Said differently, exposure to market based measures of liquidity appear to be associated with lower cross-sectional hedge fund returns. This highlights an important distinction between our measures of illiquidity and measures of liquidity based on market-liquidity betas; funds with high exposure to market liquidity are likely liquid funds with significant investments in equities and other assets whose values covary with the time-series of liquidity. The final three factors are "trend-following" factors, which are intended to capture systematic risks in Managed Futures strategies and appear to hold little significance for explaining the cross-section of returns of large funds.

Table 6 shows the correlations between $\widehat{\beta}_i$ from the FH+PS first-stage regressions, and $\overline{\log}(L_p)$ and $\overline{\log}(L_s)$. The largest correlations are with the bond-fund factor (β_{bd_mf}) at 24% and 33%, respectively. The largest negative correlations are with the equity market factor, at -21% and -12%, respectively. The correlations between our measures of illiquidity and the loadings on the Pastor-Stambaugh liquidity factors are small (and in some cases negative). The low correlations in Table 6 offer evidence that the illiquidity measures we construct are unrelated to the risk factors in standard hedge fund factor models.

Next, we look at the relationship between estimated alphas and hedge fund illiquidity. Figure 4 plots two sets of estimated alphas: those estimated as the intercept in (9) ($\beta_{0,i}^{FH+PS}$) and those estimated from the second-stage regressions (α_i^{FH+PS}), against the log of average portfolio illiquidity ($\overline{\log}(L_p)$) and investor share illiquidity ($\overline{\log}(L_s)$). In each panel funds are sorted into ten groups based on $\overline{\log}(L_p)$ or $\overline{\log}(L_s)$, respectively. The alpha values plotted on the vertical axis are averages within each group.

The results in Figure 4 show the alphas estimated from standard factor models demonstrate a clear and strong relationship with hedge fund illiquidity. Further, the return differences implied by Figure 4 between the most and least liquid portfolios are economically large, ranging from 20-50 basis-points *per month*, or 2.4-6.0% per year. For comparison, the median raw monthly return in our sample is 50 basis points. Figure 4 offers compelling evidence that portfolio and investor share illiquidity are important determinants of hedge fund returns.

While Figure 4 offers visual evidence that hedge fund illiquidity is an important omitted source of hedge funds' risk-adjusted returns, we formally test whether we can reject the null that the estimated alphas are jointly zero using the test statistic from Gibbons, Ross, and Shanken (1989). First, we sort funds into five portfolios based on the level of $\log(L_{i,s,t})$ in each period t. We then calculate mean returns within each portfolio for each period, and re-estimate equation (9) separately for each portfolio using the mean portfolio returns as the dependent variable. We then calculate:

$$\frac{T-N-K}{N}\left(1+\bar{f}'\hat{\Omega}^{-1}\bar{f}\right)^{-1}\hat{\alpha}'\hat{\Sigma}^{-1}\hat{\alpha}\sim F_{N,T-N-K}, \tag{13}$$

where N is the number of portfolios (5), K is number of factors (10), $\bar{f}$ is the vector of time series means of factors, $\hat{\Omega}$ is the covariance matrix of factors, $\hat{\alpha}$ is the vector of estimated alphas, and $\hat{\Sigma}$ is the covariance matrix of residuals. The computed F-statistic from this test is greater than 33, which implies a p-value of zero to the 13$^{\text{th}}$ decimal place. This result confirms that hedge fund alphas estimated from standard factor models are not jointly zero, and that illiquidity may be an important explanatory factor of hedge fund returns.

5.3 Hedge Fund Returns and Illiquidity

The preceding section offers evidence that hedge fund illiquidity has explanatory power for hedge funds' risk-adjusted returns. To formally test the extent to which illiquidity can explain the pricing errors estimated from the baseline two-stage regressions, we re-estimate (10) but include hedge fund portfolio or investor share illiquidity as an explanatory fund characteristic:

$$R_{i,t}=\lambda_t\log(L_{i,\ell,t})+\gamma_t'\beta_i+(a_{i,t}^{\ell}), \tag{14}$$

$$E[R_{i,t}]=\widehat{\lambda}E_i[\log(L_{i,\ell,t})]+\widehat{\gamma}'\beta_i+\widehat{\alpha}_i^{\ell}, \tag{15}$$

$$\widehat{\gamma}=\sum_t\widehat{\gamma}_t/T,\ \ \widehat{\alpha}_i^{\ell}=\sum_i\widehat{\alpha}_{i,t}^{\ell}/T,\ \ \widehat{\lambda}=\sum_t\widehat{\lambda}_t/T, \tag{16}$$

where $\ell=p,s$ denotes either log portfolio illiquidity (p) or log investor share illiquidity (s), depending on which model is estimated, and we distinguish the pricing errors in (14) from those estimated in (10) by adding the superscript $\ell=p,s$.

If hedge fund illiquidity is an important component of hedge fund returns, we should see that alphas estimated from equation (14), α_i^{ℓ}, are considerably smaller than alphas estimated from equation (10), α_i^{FH+PS}. We should also see that the absolute values of estimated alphas (the mean absolute pricing errors) are smaller as well. Table 7 shows the distribution of alphas and the absolute values of alphas estimated from our three standard factors models (FF, FH, and FH+PS) with and without the inclusion of illiquidity as an additional risk factor. The results in Table 7 confirm that illiquidity is integral component of hedge fund returns. In the FF model, investor share illiquidity reduces average alpha from 0.33% per month to 0.03% per month; investor share illiquidity can explain virtually the entirety of average alpha estimated from the Fama-French-

Carhart model. The mean absolute pricing error in the FF model is reduced from 0.41% per month to 0.29% per month, a reduction of roughly 29%. Thus, while investor share illiquidity reduces average FF alpha to nearly zero, the average absolute pricing error is still substantial. Interestingly, the reductions in alphas and mean absolute pricing errors due to portfolio illiquidity are smaller than the reductions due to investor share illiquidity, with a decrease to 0.13% and 0.31% per month for average alpha and mean absolute errors, respectively.

In the FH model, investor share illiquidity reduces average alpha from 0.37% to 0.13% (a 65% decrease) and mean pricing errors from 0.44% to 0.30% (a 32% decrease). Once again, investor share illiquidity does a better job explaining pricing errors than portfolio illiquidity, which reduces average alphas and mean pricing errors to 0.23% and 0.34%, respectively. The results from the FH+PS model are highly similar, with alphas and mean pricing errors decreasing from 0.39% to 0.14% (investor share illiquidity) and 0.24% (portfolio illiquidity), and from 0.44% to 0.30% (investor share illiquidity) and 0.35% (portfolio illiquidity), respectively. In the case of the FH and FH+PS models, our measure of investor share illiquidity appears to be roughly twice as successful at explaining average hedge fund alphas as our measure of portfolio illiquidity. We expound on this result in the following section.

Table 8 shows the difference in the effect of adding illiquidity to the FH+PS model with gross-of-fee and net-of-fee returns on the left-hand side of equation (14), respectively. While gross alphas are substantially higher than net alphas, the percentage reduction in alphas from adding investor share illiquidity or portfolio illiquidity as a fund characteristic are highly similar between gross and net returns.

One striking feature of the reduction in alphas and mean absolute pricing errors is the uniformity of reduction magnitudes across the distribution of returns. For each of the FF, FH, and FH+PS models, the amount by which alpha and the pricing errors decrease is quantitatively similar for the 10$^{\text{th}}$ through the 90$^{\text{th}}$ percentiles. This provides further evidence that the log of illiquidity is linearly related to pricing errors from standard factor model specifications.

By comparison, the bottom two rows of Table 7 show the reduction in estimated alphas and mean absolute pricing errors using the standard return-based measure of illiquidity from Getmansky, Lo, and Makarov (2004) rather than our measures of portfolio and investor share illiquidity. We use the parameter θ_0, which reflects the fraction of the true economic return comprised by the contemporaneous reported return. The inclusion of θ_0 reduces estimated FH+PS alpha by roughly 15%, and reduces the mean pricing error by 9%. Each is dramatically lower than the reductions associated with the inclusion of $\log(L_{i,s,t})$ or $\log(L_{i,p,t})$. This highlights the importance of direct measures of illiquidity for explaining the cross-section of hedge fund returns, as approximate measures based on the serial-correlation in returns offer little explanatory power in our data.

Our results contrast with Aragon (2007), who finds that hedge fund share restrictions can explain the entirety of average hedge fund alpha. We find that, while investor share restrictions can explain a substantial portion of hedge fund alpha, a sizable unexplained component remains. There are many possibilities for the difference in conclusions between these two studies. First, date ranges differ considerably, with Aragon (2007) using data from 1994-2004 and our study using data from 2013-2017. Second, our data comprise large hedge funds, whereas as the TASS data used in Aragon (2007) is skewed toward smaller funds. It may be the case that more skilled managers are better able to attract higher capital inflows from investors and therefore end up larger than less skilled funds. In this case, samples with larger funds will exhibit larger estimated alphas than samples with smaller funds.[24] Finally, it could be that the funds that choose not to report to public databases are those with bespoke strategies that carry the greatest risk of reverse engineering. If such strategies are associated with greater risk-adjusted returns, alphas estimated from non-public data may exceed those estimated from public data.

5.4 Investor Share Illiquidity vs. Portfolio Illiquidity

Section 5.3 showed that the log of investor share illiquidity, $\log(L_s)$, explained significantly more of the variation in hedge fund returns than the log of expected portfolio illiquidity, $\log(L_p)$. This result is somewhat surprising. Aragon (2007) showed that investor share restrictions could explain the entirety of average estimated alpha, but in that paper investor share restrictions were used to proxy for the illiquidity of the assets (which could not be observed). Instead, the results in Section 5.3 indicate that even when the distribution of (expected) asset liquidation horizons is observed, investor share illiquidity is still better at explaining the cross-section of hedge fund returns.

We analyze one of potentially many explanations for this result. The hypothesis is that, while *average* portfolio illiquidity earns a return premium, it does not comprise the entirety of illiquidity risk faced by the manager. For example, the second moment of L_p, which we do not observe, may command a return premium as well. This would be the case if managers not only cared about the average time needed to liquidate particular assets, but also about the liquidation time needed in worst-case scenarios or periods of asset-class or market stress. The average liquidity profile of assets may simply not well represent this risk, whereas the endogenously chosen investor share illiquidity may. Portfolio managers must take into account all risks associated with deeply illiquid assets, and are likely to set share restrictions to be consistent with

[24] In principle, the opposite could also be true. If skilled managers attract greater fund inflows, investment opportunities may diminish and returns may suffer. In this case, smaller funds may exhibit greater alphas as their fund size has yet to reach decreasing returns to scale (Berk and Green (2004)).

these risks. This hypothesis is consistent with results from Section 4.3, that shows investor share illiquidity increases more quickly than one for one with increases in portfolio illiquidity.

Consider the variance of the liquidation horizon, which can be written as:

$$Var(L_p) = Var(E[L_p]) + E[Var(L_p)]. \tag{17}$$

With our data, we can estimate only $Var(E[L_p]) = Var(\overline{L}_p)$ in equation (17). If $E[Var(L_p)]$ is also associated with a return premium, then L_p will not explain the entirety of illiquidity risk and will therefore fail to explain all of average, estimated alpha, even if illiquidity-adjusted alphas are truly zero. Yet, because hedge fund managers will optimally consider the full spectrum of illiquidity risks, including not only average liquidation periods but their potential variation as well, we should expect the investor share restrictions set by funds to more accurately reflect the totality of illiquidity risk.

Without additional data, it is difficult to directly test this hypothesis. However, the illiquidity of the financing horizon offers some independent evidence. Similarly to the questions related to portfolio and investor share illiquidity, Form PF asks funds to report the liquidity of their borrowing and available but unused cash financing.[25] We construct a measure of financing illiquidity, L_f, analogously to our construction of L_p and L_s. The hypothesis is that funds that maintain a larger mismatch between the underlying portfolio illiquidity and the illiquidity of their investors' shares are more concerned with higher-order moments of the liquidation horizon. These funds should then also opt for longer financing terms to mitigate the risk of forced asset sales from an inability to rollover debt. Thus, if funds that impose a larger mismatch between investor share illiquidity and portfolio illiquidity do so because of unobserved illiquidity risk, such funds should also choose longer financing horizons.

In Table 9, we construct a measure of "share-portfolio mismatch", measured simply as $\overline{L}_s - \overline{L}_p$, and evaluate its association with the length of the financing horizon. Columns (1) and (2) of Table 9 shows that funds that have a larger share-portfolio mismatch also choose longer average financing horizons. This suggests that funds that extend the duration of one source of funding (investor capital), also extent the duration of the other source of funding (borrowing), holding portfolio illiquidity constant. This offers preliminary evidence that unobserved illiquidity risks may be better represented by investor share illiquidity rather than expected portfolio illiquidity.

[25] Specifically, Form PF asks funds to "Divide [the aggregate dollar amount of borrowing by and cash financing available to the reporting fund (including all drawn and undrawn, committed and uncommitted lines of credit as well as any term financing)] among the periods specified below depending on the longest period for which the creditor is contractually committed to provide such financing." The periods specified are the same as for the portfolio and investor share illiquidity questions: 0-1 days, 2-7 days, 8-30 days, 31-90 days, 91-180 days, 181-365 days, and greater than 365 days.

Alternatively, the relationship between the share-portfolio mismatch and financing illiquidity may be driven by funds with strategies that produce uncertain payoff horizons. Funds that are unsure when a particular bet may pay off are likely to set tighter share restrictions to avoid being forced to unwind a position early due to redemption requests. Evidence of this is found in Giannetti and Kahraman (2017), who show that open-ended mutual funds and hedge funds with looser share restrictions are less likely to buy mispriced assets when the time to correction is less predictable. We test this mechanism with a dummy variable for whether the fund follows an event-driven strategy. Such strategies are often based on trades in highly-liquid equity and derivatives markets, but are also likely to have rather uncertain payoff dates. If funds set tighter share restrictions due to uncertain payoff horizons, we should see a positive association between the share-portfolio mismatch and the event-driven dummy.

First, column (3) of Table 9 demonstrates that event-driven funds have roughly average reported portfolio illiquidity as measured by $\overline{L}_p$. Columns (4) and (5) show that, with and without controls, the event driven dummy has an economically large coefficient that is statistically significant at the 5% level. This offers some evidence that funds also account for uncertainty in the investment horizon when setting share restrictions, consistent with Giannetti and Kahraman (2017). In column (6), we include both $\overline{L}_f$ and the event-driven dummy in the regression simultaneously. The coefficients on each remain statistically significant and change in magnitude only slightly. This suggests that both the uncertainty of the payoff horizon and higher-order moments of the liquidation horizon are important components of the share-portfolio mismatch.

5.5 The Illiquidity Return Premium

The parameter $\widehat{\lambda}$ in equation (14) represents the illiquidity return premium. Specifically, the parameter $\widehat{\lambda}$ gives the additional, *ceteris paribus* expected return earned from an additional log-day of illiquidity — associated either with portfolio illiquidity or investor share illiquidity depending on the specification. This gives the return premium estimated in this paper a natural economic interpretation, unlike those estimated from portfolio sorts based on return serial correlation or other measures of illiquidity based on returns (Getmansky, Lo, and Makarov (2004)). Based on the evidence from sections 5.3 and 5.4, in what follows we use investor share illiquidity, $\overline{\log}(L_s)$, as our measure of illiquidity, but interpret it as a proxy for the true underlying illiquidity of the fund's assets.

The two-stage Fama-Macbeth procedure estimated in equation (14) provides coefficient estimates of $\widehat{\lambda}_{gross} = 0.0681$ and $\widehat{\lambda}_{net} = 0.513$ for the illiquidity premium estimated from gross- and net-of-fee returns, respectively. The return premium estimated from gross returns provides an estimate of the compensation earned by the underlying assets held by the fund; it is a measure that is comparable across markets and market-participants. The return premium estimated from net returns is the additional compensation earned by the hedge fund investor from an increase in the illiquidity of fund assets, after accounting for the dilution of returns due to investments made through an intermediary. Figure 5 plots the expected, *ceteris paribus* returns implied by these estimates of the illiquidity premium. As shown in Figure 5, the return premiums estimated from gross and net returns are largely similar. Further, the t-statistic on $\widehat{\lambda}$ estimated from net returns is 7.52, which allows us to strongly reject the null that the return premium earned by fund investors is zero. Taken together, this indicates that a substantial portion of the illiquidity premium earned by fund assets are passed through to fund investors.

The parameter $\widehat{\lambda}$ implies that a fund in the 5^{th} percentile of illiquidity will, all else equal, earn just over 20 basis points per month on average across all strategies, of which roughly 19 basis points are passed to fund investors. A fund in the 95^{th} percentile will earn over 55 basis points a month of which 45 are passed to the investor. This gives a gross illiquidity premium of nearly 35 basis points per month, equivalent to an annual return premium of nearly 4.3% per year. Investors earn an illiquidity premium of around 23 basis points per month between the 5^{th} and 95^{th} illiquidity-percentile funds, or around 2.8% per year. Stated differently, $\widehat{\lambda}_{gross}$ implies that an additional log-day of illiquidity is associated with roughly an additional 0.068 basis points of returns per month, or roughly 82 basis points per year.

The ratio $\widehat{\lambda}_{net}/\widehat{\lambda}_{gross}$ provides the fraction of the illiquidity premium passed through to fund investors. The coefficient estimates imply that roughly 75% of the illiquidity premium is ultimately captured by investors. To our knowledge, no such estimate exists elsewhere in the literature, because it requires information on funds' asset illiquidity as well as measures of gross and net returns.

To demonstrate the importance of more direct measures of hedge fund illiquidity, such as $\log(L_s)$ and $\log(L_p)$, we contrast $\widehat{\lambda}$ estimated from equation (14) with the return premium estimated by using the parameter θ_0 from Getmansky, Lo, and Makarov (2004) as the measure of liquidity. Not only are we unable to reject the null that the illiquidity return premium based on θ_0 is greater than zero, but the return premium we estimate is actually negative; more liquid funds earn a premium based on liquidity as measured by θ_0. These parameter estimates indicate that the return premiums estimated from a standard measure of hedge fund illiquidity based on the serial-correlation in fund returns are likely to be problematic and difficult to interpret. We leave a more thorough treatment of this issue for future research.

5.6 Construction of an Illiquidity Risk Factor and Return Smoothing

The previous sections have established that illiquid funds earn considerably higher returns on average than liquid funds, and that fund illiquidity can explain between 38-64% of alphas estimated from standard factor models. However, to this point we have yet to address whether these higher returns are compensation for exposure to a systematic illiquidity risk factor.

Below, we construct a time-series risk factor based on our measures of illiquidity. We find that such a risk factor has no predictive power for the cross-section of fund returns, which suggests it is the level of illiquidity rather than systematic illiquidity risk that drives our results. This finding is consistent with Bongaerts, de Jong, and Driessen (2017), who find that the illiquidity level contributes around 1% of the 1.9% excess bond return, whereas equity market liquidity risk, equity market risk, and volatility risk account for only 0.3%.

5.6.1 The Illiquidity Factor and Return Smoothing

Specifically, we sort funds into five groups based on their average investor share illiquidity, $\overline{L}_s$. Because our sample includes just under five years of returns, we sort funds only once. This is likely not too problematic because investor share illiquidity is highly persistent. In each month and year, we construct our factor as the mean return to the 20% most illiquid funds minus the mean return to the 20% most liquid funds. We note that this factor does not constitute a traded portfolio, because hedge fund shares are not securities. However, the assets held by the funds are traded.

The top panel of Figure 6 shows the time series distribution of the illiquidity risk factor. The mean value of the factor is 0.12% with a standard deviation of 0.71%. On average, the 20% most illiquid funds outperform the 20% most liquid by roughly 1.5% per year. For comparison, the top panel also includes the traded liquidity factor from Pastor and Stambaugh (2003), which has a correlation of 31% with the illiquidity risk factor constructed here. This suggests the illiquidity factor does capture some variation in market-wide liquidity.

The bottom panel of Figure 6 shows the time series returns for each of the two portfolios that comprise the illiquidity factor: the average returns to the 20% most illiquid and liquid funds, respectively. The bottom panel offers a few interesting takeaways. The first is that the two portfolios appear to move rather tightly together over time; indeed, their correlation is nearly 81% over the sample period. The second is that the average return to the most liquid funds is considerably more volatile than the most illiquid funds. The standard deviation of the liquid portfolio is 1.15% per month, compared to 0.726% for the illiquid portfolio.

The lower volatility of average returns associated with illiquid funds highlights an important consideration about the illiquidity premium and *return smoothing*. A well-known property of hedge fund returns is that stale prices due to illiquid assets give rise to autocorrelation in fund returns (Getmansky, Lo, and Makarov (2004)). In the case of smoothed returns, betas estimated from contemporaneous risk factors will understate the true exposure to the factor, and lower betas mean higher alphas all else equal (Asness, Krail, and Liew (2001)).

We examine the extent of return smoothing for each of the five illiquidity-sorted portfolios in Table 10. Column one lists the five portfolios from most to least liquid, and column two shows the average investor share illiquidity within each portfolio. In column three, we calculate the first-order autocorrelation coefficient based on the average monthly returns to each portfolio. Indeed, consistent with Khandani and Lo (2011), column three shows that the autocorrelation coefficient for the average returns to the illiquid portfolio (0.293) is considerably higher than for the liquid portfolio (0.012). Column four calculates the autocorrelation coefficient for monthly returns at the fund level, then averages those values within each portfolio. The coefficients estimated for the first three portfolios are relatively similar across the two measures, but the average fund level coefficients are considerably lower than the portfolio coefficients for the two most illiquid portfolios. Even so, column four shows a sizeable difference between the average autocorrelations in the first and fifth illiquidity portfolios.

However, Table 10 also shows that it is unlikely the higher values of alpha estimated for illiquid portfolios are the result of return smoothing. First, the illiquidity premium shows up even in raw returns. Column five shows the most illiquid portfolio has an average monthly return that is roughly 23 basis points higher than the most liquid portfolio, which compounds to an average of 2.8% higher annual returns, and has a two-sample t-statistic of 3.56, suggesting the average return difference is statistically different from zero. Next, the differences in average alphas seem largely unrelated to the autocorrelation coefficients. The largest differences in estimated alphas arise from comparisons of portfolios one and two, and four and five, with risk-adjusted return differences of 0.13% and 0.20%, respectively. However, the differences in autocorrelations between these two sets of portfolios are small. The difference between the first and second portfolio is 0.031 (.043 versus .012) or 0.048 (.060 versus .012), depending on the method, and for the fourth and fifth portfolios is 0.019 (.293 versus .275) or 0.0126 (0.136 versus 0.110). The largest difference in autocorrelations comes from a comparison of portfolios three and four (0.169, column three), or two and three (0.061, column four), but the difference in estimated alphas is smallest for this group (0.09%). Thus, it appears unlikely that the autocorrelation in returns is the primary characteristic generating observed differences in either the raw returns or estimated alphas.

Finally, while the lower volatility (not smoothing) of returns earned by the most illiquid funds will also shrink the estimated betas, this likely reflects appropriate economic correlations. Funds with tighter share restrictions can more efficiently manage the sale of assets, taking longer to find buyers willing to pay higher prices and selling large blocks of securities over longer periods. In this case, we should expect the returns to illiquid funds to be less volatile, and to covary less with traded risk factors. The increased alpha that arises from this type of beta-shrinkage therefore does represent higher risk-adjusted economic returns.

5.6.2 The Illiquidity Factor and the Cross-Section of Returns

To examine whether it is exposure to the illiquidity risk factor that in part generates higher expected returns, we repeat the Fama-MacBeth two-stage regressions outlined in section 5.1 but include the illiquidity risk factor in the first stage regression. If exposure to the illiquidity risk factor explains hedge fund returns, the beta coefficients estimated from the first stage should produce an economically large and statistically significant coefficient in the second stage, offering evidence of priced illiquidity risk.

We find no such relationship in our data. The second stage coefficient on the illiquidity factor betas range from -1.80 to 1.48 over the 57 monthly cross-sectional regressions, with a mean of -.0075 and a standard deviation of 0.7255. This results in a t-statistic of -0.07, far from statistically significant. Although expected raw returns covary slightly with illiquidity factor betas estimated from the first stage, the economic significance is small, and the second-stage coefficients demonstrate that once other risk factors are included no explanatory power remains in the illiquidity risk factor.

In summary, while the constituent portfolios that comprise the illiquidity factor offer important insights into the volatility and smoothness of hedge fund returns, as well as their implications for estimated alphas, we find no statistically significant relationship between exposure to the illiquidity risk factor and the cross-section of fund returns. Our results indicate that the entirety of the illiquidity premium we estimate originates from the expected transactions costs that arise from trading illiquid securities, such as price impact, or from compensation for risk that the undiversified hedge fund specialists requires to be a marginal investor in illiquid markets.

6 Alphas and Fees

Section 4.2 showed that portfolio illiquidity and investor share illiquidity are highly correlated. This is consistent with managers passing some of the illiquidity risk associated with portfolio assets to fund investors through the imposition of share restrictions, such as lockups and redemption notice periods. Further, Section 5.5 showed that the illiquidity premium estimated from gross returns is 25% higher than the premium earned by fund investors. This means that fund managers are capturing some of the illiquidity premium, perhaps because not all of the risks managers face by investing in illiquid assets are passed to fund investors.

If the fees charged by funds with more illiquid assets are the same as those charged by funds with more liquid assets, the fraction of the illiquidity premium captured by funds is not particularly interesting. Suppose that all funds charge "2 and 20", a 2% management fee and a 20% performance fee; in this case the fraction of the illiquidity premium retained by managers is 20% by construction (20% of *all* sources of additional returns are captured by the manager). For the fraction of the illiquidity premium captured by managers to be interesting, it must be that the fees funds charge are a function of illiquidity (this is the assumption in the reduced-form asset pricing model described in Section 2). If funds require additional compensation for undiversifiable illiquidity risk, it must be that funds charge higher *rates* of fees in response to larger investments in illiquid assets. This compensation could be either because managers choose to bear some of the illiquidity risk and therefore share in the illiquidity premium, or because funding illiquidity risks that result from borrowing, such as margin risk, cannot be fully passed to investors.

The testable hypothesis is that investor fees should be related to the amount of illiquidity in the fund's portfolio. To establish facts, we first examine whether average realized fees — calculated as the mean of the annual reported gross rate of return minus the annual net rate of return — are partially explained by either investor share illiquidity or manager skill.[26] We emphasize that this analysis examines *realized* fees, which comprise the individual fee components such as the management fee rate (for example, 1.5% of assets under management), the incentive fee rate (for example, 20% of returns), and the imposition of a high water mark or hurdle rate.

Figure 7 plots average realized fees versus $\alpha_i^{s,g}$, and versus $\overline{\log}(L_{i,s})$ after controlling for $\alpha_i^{s,g}$. We add the superscript g to the alpha variable to emphasize that these are alphas estimated from gross of fee returns. This is important because Berk and Green (2004) show that in a competitive market for fund managers, gross of fee alphas should be correlated with fees but net of fee alphas should be the same across funds with different levels of skill. The left panel shows that funds with higher estimated alphas indeed earn higher average fees. The right panel shows that funds with more illiquid assets also earn higher average fees *after controlling for* $\alpha_i^{s,g}$. This offers consistent evidence that managers of illiquid assets do not pass all of the undiversifiable illiquidity risk to fund investors, and instead require some additional compensation for the illiquidity risk exposure of their personal wealth invested in the fund.

To better understand the relationships in Figure 7, we first examine whether illiquidity mechanically leads to higher fees. Most hedge funds have fee structures that comprise two components, a management fee equal to some percentage of the total assets under management (often 1.5% annually), and an incentive fee equal to some percentage of the returns earned in that period (often 20%). Funds that have positive returns will therefore earn both the management fee and the incentive fee, whereas funds with negative

[26] Form PF provides little guidance on how funds should report gross and net fees. The form states only that funds should "Provide the reporting fund's gross and net performance, as reported to current and prospective investors (or, if calculated for other purposes but not reported to investors, as so calculated)"

performance will only earn the management fee. Similarly, many funds employ a "high-water mark", which stipulates that incentive fees can only be charged on returns earned above the previous high value of NAV, and/or a "hurdle rate", that specifies a minimum rate of return below which the incentive fee is not paid.[27] For example, a fund that employs a high-water mark and earns a 10% return in January, establishing a new high-value for NAV, but then loses 6% in February, will only earn the incentive fee in March on returns that exceed the 6% the fund lost in February. Funds that are well below the high-water mark may go many months without earning the incentive fee even after earning positive returns. Thus, funds with higher returns due to the illiquidity premium may earn mechanically higher average fees because they are more likely to earn the incentive fee, are less likely to fall below the high-water mark, and are more likely to exceed the hurdle rate.

Table 11 examines the relationship between illiquidity and fees more formally. Column (1) shows that, after controlling for size (net assets) and strategy, higher alpha funds based on gross returns indeed earn higher fees. Column (2) shows that this relationship remains after controlling for average investor portfolio illiquidity; further, portfolio illiquidity predicts higher average fees even after controlling for estimated alpha. Column (3) includes the proportion of months in which the fund has positive (gross) returns as an attempt to directly control for the likelihood of earning the incentive fee or exceeding the high-water mark. The fraction of positive-return months does explain some of average realized fees, and reduces the coefficient on estimated alpha by roughly 40%, which remains highly statistically significant. Further, the coefficient on portfolio illiquidity is reduced from 0.159 to 0.112, but also remains highly statistically significant.

However, the results in the previous sections suggest that investor share illiquidity may be a better measure of the total portfolio illiquidity than the average illiquidity of the assets. In column (4) of Table 11, we include the log of investor share illiquidity in addition to the log of portfolio illiquidity. Consistent with our earlier findings, average fees are strongly related to investor share illiquidity but are unrelated to portfolio illiquidity once investor share illiquidity is included. The coefficient estimate suggests a one log-day increase in investor share illiquidity is associated with a 0.23% larger annual realized fee, and is highly statistically significant.

This result may seem paradoxical. Taken literally, the results from column (4) suggest that funds that impose tighter share restrictions also charge higher fees, *controlling for manager skill*; for an investor, this would be a lose-lose proposition. However, placed in the context of unobserved illiquidity risk, these results are sensible. As the illiquidity of the fund's assets increases, the manager passes some of the costs to

[27] High-water marks are imposed on an investor-by-investor basis; that is, based on the highest previous NAV for each particular investor, which may vary by investors based on when the investor first invested in the fund.

investors in the form of tighter share restrictions, and maintains some of the risk himself, requiring higher fees as compensation. As asset illiquidity grows, so too will fees and investor share illiquidity, as the fund manager and fund investors are each sharing a portion of the illiquidity costs (and therefore the illiquidity premium). Thus, true unobserved illiquidity leads to a positive correlation between investor share illiquidity and fees.

In column (5), we offer additional evidence of this interpretation. One way in which managers will not pass the entirety of the illiquidity risks to fund investors is if some portion of those risks arise from financing liquidity risks that result from leveraged positions. In this case, we should see more leveraged funds charge higher fees *ceteris paribus*, as higher leverage implies higher illiquidity risks that cannot be passed to investors through share restrictions. Column (5) finds evidence of such an effect. Controlling for estimated skill, the fraction of positive-return months, portfolio illiquidity, and investor share illiquidity, the coefficient on average fund leverage is statistically significant and economically large. A fund that is leveraged two-to-one will earn a 0.067% higher annual fee on average than a similar fund with no leverage. Given the host of additional controls in this specification, the results in column (5) appear to be strong independent evidence that the additional fees associated with illiquidity arise as explicit compensation for the undiversifiable risks faced by managers of illiquid assets.

Next, we attempt to decompose fees into their constituent parts: a management fee based on a constant proportion of assets under management, and an incentive fee based on a proportion of fund returns. Unfortunately, we do observe these fees directly, nor do we observe whether the fund is above the high-water mark (or whether the fund even observes a high-water mark). First, we restrict our sample to observations for which the gross return is greater than or equal to the net return (this excludes data errors and funds offering rebates). We estimate the management fee by taking the maximum fee charged by the fund during months in which it has a negative return; this should eliminate the incentive fee and leave only the management fee as the fee earned in these months. We then winsorize at the 5% and 95% levels. To estimate the incentive fee, which is often charged either based on quarterly or annual performance, we divide the total fee minus the estimated management fee by the total gross return minus the management fee, all in annual percentages. We then winsorize the estimated incentive fee at the 10% and 90% levels due to rather large outliers. This procedure produces a median incentive fee estimate of 17.84% with a mean of 15.10%, and a median management fee of 1.21% with a mean of 1.35%. Each is consistent with industry standards of 15-20% incentive fees and 1.0-2.0% management fees.

To test the credibility of our fee estimates, we use a sample of matched funds from the Form PF and Lipper TASS databases from Barth and Wermers (2018). There are 142 qualifying hedge funds in the Form

PF data that also report to TASS. The TASS database contains explicit information on both the management and incentive fees charged by funds, and therefore allows us to directly test our fee estimates for the subset of funds that appear in both data sets. The vast majority of matched funds report to TASS an incentive fee of 20%; of those funds, the mean incentive fee we estimate in Form PF data is 19.23% with a median of 19.84% and an interquartile range of 17.32-22.09%. The vast majority of matched funds report to TASS management fees of either 1.5% or 2.0%. Of those that charge 1.5%, we estimate a management fee with an interquartile range of 0.96-1.7%, with a mean of 1.48% and a median of 1.57%. For those that charge a 2% management fee, we estimate an interquartile range of 1.27-2.06%, with a mean of 1.75% and a median of 2.06%. In total, we view these results as strongly supportive of the credibility of our fee estimates.

To establish facts, in column (6) of Table 11 we regress the estimated incentive fee on the estimated management fee. The coefficient is positive, large, and highly statistically significant. This suggests that managers who charge larger incentive fees also charge higher management fees, perhaps reflecting that skilled managers can charge higher fees along both dimensions.

In column (7) of Table 11, we regress the estimated incentive fee on estimated gross-of-fee alpha, average investor share illiquidity, average leverage, and other controls. The estimated incentive fee is economically related to both estimated alpha and investor share illiquidity. The coefficients on each are also highly statistically significant. A ten basis point higher monthly alpha is associated with a 0.3% higher incentive fee. An extra log-day of investor share illiquidity is associated with an additional 1.98% higher incentive fee after controlling for gross-of-fee alpha. Average leverage is also economically large and statistically significant, with a fund leveraged two-to-one on average charging roughly 30 basis points higher incentive fees than an unleveraged fund.

Alternatively, column (8) shows that while higher alpha is associated with a higher management fee, investor share illiquidity appears uncorrelated with the management fee. This is consistent with models of intermediary-based asset pricing (He and Krishnamurthy (2012)), and our interpretation of the relationship between illiquidity and fees; because the management fee is fixed in each period, only the incentive fee will reward managers who are more highly invested in illiquid assets with a higher share of the illiquidity premium. Consistent with this interpretation, leverage is also unrelated to the estimated management fee, suggesting that compensation for financing liquidity risk is also paid through the incentive fee.

Because our measures of the incentive and management fees are estimated with noise, we test our results using the Lipper TASS database. We find qualitatively similar results, although estimated magnitudes differ. Generally, we find that in the TASS data the incentive fee is positively correlated with the management fee, and that funds that have tighter share restrictions as measured by lockups and redemption notice periods charge higher incentive fees, but do not charge higher management fees (in fact, we find a negative relationship between management fees and investor share restrictions in TASS). In summary, the results in Table 11 appear to hold qualitatively in the TASS data as well.

We also note that the inclusion of gross-of-fee alphas in Table 11 allows us to rule out one possible alternative hypothesis relating investor share restrictions to higher fees. If more skilled managers have greater bargaining power relative to investors, they could simultaneously charge higher fees and impose longer lock-ups. Indeed, gross alpha is highly related to both the incentive and management fee. However, we continue to find a strong relationship between investor share illiquidity and the incentive fee even after controlling for manager skill. Leverage also remains associated with the incentive fee. This suggests that bargaining power is insufficient to explain the positive correlation between investor share illiquidity and fees.

Finally, based on the model of Berk and Green (2004), we examine whether we find similar results with net-of-fee alphas. Table 12 repeats the specifications in Table 11, but includes net-of-fee alpha instead of gross-of-fee alpha. We find that once the fraction of months with positive returns is included in the regression, net-of-fee alpha is no longer associated with higher fees. We similarly find no relationship between net-of-fee alpha and the management or incentive fees. This suggests that all of the additional returns produced by more skilled managers are captured by the manager through higher fees. This result is consistent with a competitive market for hedge fund managers. However, despite the lower correlations between net-of-fee alpha and fees, the coefficients on investor share illiquidity are virtually unchanged from Table 11 in each specification. This provides another robustness check for the associations between investor share illiquidity and fees.

7 Conclusion

We use data from the first systematic, regulatory collection of data on large hedge funds to investigate a number of empirical and economically meaningful relationships related to asset illiquidity in the asset management industry. We argue that the potential price impact associated with forced sales of illiquid assets constitutes an undiversifiable risk for hedge fund managers. In response, managers may restrict investors shares and charge higher fees.

We first document that investor share illiquidity and portfolio illiquidity are strongly positively correlated. This suggests at least some of the costs of illiquid assets are passed to fund investors. For hedge fund returns, we find that between 38-64% of hedge fund alpha can be explained by illiquidity. Surprisingly, investor share illiquidity has much greater explanatory power for risk-adjusted returns than portfolio illiquidity. We also estimate a sizeable illiquidity premium; the *ceteris paribus* return differential between funds in the 5^{th} and 95^{th} percentiles of illiquidity is nearly 4.3% per year. Premiums estimated from alternative measures of illiquidity, such as return smoothing, fail to produce statistically significant results. We further find that roughly 75% of the illiquidity premium is passed to fund investors, implying that 25% of the illiquidity premium is retained by the manager.

Consistent with this finding, we show that hedge fund fees are higher for illiquid funds, even after controlling for manager skill and a host of other possible explanatory variables. Consistent with intermediary-based asset pricing theories, all of the association between illiquidity and fees arises through the performance fee. Finally, we find that while gross-of-fee alphas are strongly related to realized fees, net-of-fee alphas are uncorrelated with fees. This suggests the entirety of excess returns due to manager skill are captured by the fund managers.

Our results provide a novel estimate of the return premium earned by illiquid assets, based on the portfolios of the marginal investors in illiquid markets rather than on asset-specific estimates of price impact or short-term reversals. Our results also have important implications for the interpretation of alpha from standard factor models as manager skill. Without controlling for the illiquidity of portfolio assets or investor shares, the unexplained part of hedge fund returns will comprise both skill and the illiquidity premium. Further, we document a new potential compensated risk in the limits to arbitrage literature — price impact from forced asset sales in intermediated asset management. Our empirical results provide support this interpretation, and our data allow us to rule out some alternative explanations. In summary, we document new economic relationships between managers and investors in the hedge fund industry.

References

ACHARYA, V. V., AND L. H. PEDERSEN (2005): "Asset pricing with liquidity risk," *Journal of Financial Economics*, 77(2), 375 – 410.

AIKEN, A. L., C. P. CLIFFORD, AND J. A. ELLIS (2015): "Hedge funds and discretionary liquidity restrictions," *The Journal of Financial Economics*, 116(1), 197–218.

AMIHUD, Y. (2002): "Illiquidity and stock returns: cross-section and time-series effects," *Journal of Financial Markets*, 5(1), 31 – 56.

AMIHUD, Y., AND H. MENDELSON (1986): "Asset pricing and the bid-ask spread," *Journal of Financial Economics*, 17(2), 223 – 249.

ARAGON, G. (2007): "Share restrictions and asset pricing: Evidence from the hedge fund industry," *Journal of Financial Economics*, 83(1), 33–58.

ARAGON, G. O., T. ERGUN, M. GETMANSKY, AND G. GIRARDI (2017a): "Hedge Fund Liquidity Management," *Securities and Exchange Commision Working Paper*.

——— (2017b): "Hedge Funds: Portfolio, Investor, and Financing Liquidity," *Securities and Exchange Commision White Paper*.

ASNESS, C. S., R. J. KRAIL, AND J. M. LIEW (2001): "Do Hedge Funds Hedge?," *The Journal of Portfolio Management*, 28(1), 6–19.

BARTH, D., AND R. WERMERS (2018): "Biases in Public Hedge Fund Data," *Working Paper*.

BERK, J. B., AND R. C. GREEN (2004): "Mutual Fund Flows and Performance in Rational Markets," *Journal of Political Economy*, 112(6), 1269–1295.

BOLLEN, N. P. B., AND V. K. POOL (2008): "Conditional Return Smoothing in the Hedge Fund Industry," *The Journal of Financial and Quantitative Analysis*, 43(2), 267–298.

BONGAERTS, D., F. DE JONG, AND J. DRIESSEN (2017): "An Asset Pricing Approach to Liquidity Effects in Corporate Bond Markets," *The Review of Financial Studies*, 30(4), 1229–1269.

CAO, C., Y. CHEN, B. LIANG, AND A. W. LO (2013): "Can hedge funds time market liquidity?," *Journal of Financial Economics*, 109(2), 493 – 516.

CARHART, M. M. (1997): "On Persistence in Mutual Fund Performance," *The Journal of Finance*, 52(1), 57–82.

COVAL, J., AND E. STAFFORD (2007): "Asset fire sales (and purchases) in equity markets," *Journal of Financial Economics*, 86(2), 479 – 512.

DING, B., H. A. SHAWKY, AND J. TIAN (2009): "Liquidity shocks, size and the relative performance of hedge fund strategies," *The Journal of Banking and Finance*, 33(5), 883–891.

FALATO, A., A. HORTASCU, D. LI, AND C. SHIN (2017): "Fire-Sale Spillovers in Debt Markets," *Working Paper*.

FAMA, E. F., AND K. R. FRENCH (1992): "The Cross-Section of Expected Stock Returns," *The Journal of Finance*, 47(2), 427–465.

FAMA, E. F., AND J. D. MACBETH (1973): "Risk, Return, and Equilibrium: Empirical Tests," *Journal of Political Economy*, 81(3), 607–636.

FLOOD, M., AND P. MONIN (2016): "Form PF and the Systemic Risk of Hedge Funds: Risk-measurement precision for option portfolios," *Journal of Alternative Investments*, 18(4), 125–147.

FLOOD, M., P. MONIN, AND M. BANDYOPADHYAY (2015): "Gauging Form PF: Data Tolerances in Regulatory Reporting on Hedge Fund Risk Exposures," Working Paper No. 15-13, Office of Financial Research.

FUNG, W., AND D. A. HSIEH (2001): "The Risk in Hedge Fund Strategies: Theory and Evidence from Trend Followers," *The Review of Financial Studies*, 14(2), 313.

——— (2004): "Hedge Fund Benchmarks: A Risk-Based Approach," *Financial Analysts Journal*, 60(5), 65–80.

GABAIX, X., A. KRISHNAMURTHY, AND O. VIGNERON (2007): "Limits of Arbitrage: Theory and Evidence from the Mortgage-Backed Securities Market," *The Journal of Finance*, 62(2), 557–595.

GETMANSKY, M., A. LO, AND I. MAKAROV (2004): "An econometric model of serial correlation and illiquidity in hedge fund returns," *Journal of Financial Economics*, 74(3), 529–609.

GIANNETTI, M., AND B. KAHRAMAN (2017): "Open-end Organizational Structures and Limits to Arbitrage," *Swedish House of Finance Research Paper No. 14-09*.

GIBBONS, M. R., S. A. ROSS, AND J. SHANKEN (1989): "A Test of the Efficiency of a Given Portfolio," *Econometrica*, 57(5), 1121–1152.

GROMB, D., AND D. VAYANOS (2002): "Equilibrium and welfare in markets with financially constrained arbitrageurs," *Journal of Financial Economics*, 66(2), 361–407.

HE, Z., AND A. KRISHNAMURTHY (2012): "A model of capital and crises," *The Review of Economic Studies*, 79(2), 735–777.

JOTIKASTHIRA, C., C. LUNBLAD, AND T. RAMADORAI (2012): "Asset Fire Sales and Purchases and the International Transmission of Funding Shocks," *The Journal of Finance*, 67(6), 2015–2050.

KHANDANI, A. E., AND A. W. LO (2011): "Illiquidity Premia in Asset Returns: An Empirical Analysis of Hedge Funds, Mutual Funds, and US Equity Portfolios," *Quarterly Journal of Finance*, 01(02), 205–264.

KYLE, A. S., AND W. XIONG (2001): "Contagion as a wealth effect," *The Journal of Finance*, 56(4), 1401–1440.

PASTOR, L., AND R. F. STAMBAUGH (2003): "Liquidity Risk and Expected Stock Returns," *Journal of Political Economy*, 111(3), 642–685.

SADKA, R. (2006): "Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk," *Journal of Financial Economics*, 80(2), 309 – 349.

——— (2010): "Liquidity risk and the cross-section of hedge-fund returns," *The Journal of Financial Economics*, 98(1), 54–71.

SHLEIFER, A., AND R. W. VISHNY (1992): "Liquidation Values and Debt Capacity: A Market Equilibrium Approach," *The Journal of Finance*, 47(4), 1343–1366.

SHLEIFER, A., AND R. W. VISHNY (1997): "The Limits of Arbitrage," *The Journal of Finance*, 52(1), 35–55.

SIRIWARDANE, E. (2018): "Limited Investment Capital and Credit Spreads," *Journal of Finance*, Forthcoming.

TEO, M. (2011): "The liquidity risk of liquid hedge funds," *The Journal of Financial Economics*, 100(1), 24–44.

8 Tables and Figures

8.1 Tables

Table 1: Summary Statistics

Variable	**Mean**	**Std. Dev.**	25th **Percentile**	50th **Percentile**	75th **Percentile**	**Obs.**
NAV ($Billions)	1.13	2.19	0.17	0.51	1.15	2,579
GAV ($Billions)	2.04	7.09	0.25	0.68	1.72	2,579
GAV /NAV	3.43	38.56	1.01	1.16	1.67	2,579
Monthly Obs./Fund	35.73	21.02	16.00	44.00	57.00	2,579
Quarterly Obs./Fund	13.26	5.92	8.00	15.00	19.00	2,579
Strategy						
Equity						807
Multi-strategy						319
Event Driven						191
Relative Value						217
Credit						226
Macro						161
Managed Futures						*
Invest in Other Funds						*
Other						539

Table 1 reports summary statistics for our main sample of 2,518 funds. Each statistic reported is the average value for the fund over the full set of observations for which we observe it. Stars are used to replace data for strategies with too few observations to meet the appropriate data aggregation criteria. These observations are still included in our other analyses, where they are properly aggregated.

Table 2: Portfolio, Investor, and Financing illiquidity

		Portfolio Liquidity ($\overline{L}_p$)			Investor Liquidity ($\overline{L}_s$)		
	Pct. Obs.	**Mean**	**Median**	**SD.**	**Mean**	**Median**	**SD.**
All	100.0	64.7	12.2	106.9	164.6	135.5	137.5
Equity	31.3	24.2	6.3	53.4	123.7	77.2	107.0
Multi-strategy	12.4	59.5	23.5	89.6	175.1	174.6	129.8
Event Driven	7.4	91.5	41.3	106.9	254.2	283.6	115.7
Relative Value	8.4	45.2	12.4	74.8	154.7	135.5	130.1
Credit	8.8	105.9	28.1	128.0	220.5	252.0	139.9
Macro	6.2	7.0	2.9	12.6	75.2	60.5	80.3
Other	20.9	135.2	55.7	147.2	217.6	304.0	159.9

Table 2 reports summary statistics for the averages of our three illiquidity measures: portfolio illiquidity ($\overline{L}_p$), investor share illiquidity ($\overline{L}_s$), and financing illiquidity ($\overline{L}_f$). This table includes only the 2,518 funds in our main sample. Stars are used to replace data for strategies with too few observations to meet the appropriate data aggregation criteria. These observations are still included in our other analyses, where they are properly aggregated.

Table 3: Measures of Liquidity

	(1)	(2)	(3)	(4)	(5)
	$\overline{L}_p$	$\overline{L}_p$	$\log \overline{L}_p$	L_p	L_p
$\overline{L}_s$	0.499***	0.453***	0.009***		
	(0.012)	(0.012)	(0.000)		
L_s				0.422***	0.057**
				(0.018)	(0.022)
Constant	-17.454***	6.769	1.429***	9.754	43.928***
	(2.521)	(5.864)	(0.103)	(9.281)	(5.292)
Strategy Dummies	No	Yes	Yes	Yes	Yes
Leverage Controls	No	Yes	Yes	Yes	Yes
Size Controls	No	Yes	Yes	Yes	Yes
Time FE	No	No	No	Yes	No
Fund FE	No	No	No	No	Yes
Adjusted R^2	0.411	0.470	0.526	0.432	0.010
Observations	2579	2579	2579	20645	20645

Table 3 reports regression results of portfolio liquidity L_p on investor share liquidity L_s, as well as their averages. Size controls include net asset value, leverage controls include the ratios of gross assets, borrowing, and gross notional exposures to net assets, and strategy dummies are based on broad strategy category. Time FE are fixed-effects by quarter and year. In column (4), standard errors are clustered on both quarter-year and fund. In column (5), standard errors are clustered only at the fund level. * denotes significance at the 10% level, ** at the 5% level, and *** at the 1% level.

Table 4: Summary Statistics: Returns and Time-Series Alphas

Strategy	**Return Measure**	**Mean**	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}	**St. Dev.**
All	Avg. $R_{i,t}$	0.52	0.01	0.28	0.50	0.78	1.08	0.46
	$\widehat{\beta}_{0,i}^{FF}$	0.20	-0.31	-0.06	0.21	0.45	0.72	0.42
	$\widehat{\beta}_{0,i}^{FH}$	0.26	-0.28	-0.01	0.27	0.53	0.78	0.43
	$\widehat{\beta}_{0,i}^{FH+PS}$	0.35	-0.21	0.06	0.36	0.63	0.91	0.47
Equity	Avg. $R_{i,t}$	0.68	0.17	0.39	0.65	0.95	1.27	0.47
	$\widehat{\beta}_{0,i}^{FF}$	0.11	-0.39	-0.16	0.10	0.38	0.61	0.43
	$\widehat{\beta}_{0,i}^{FH}$	0.13	-0.38	-0.18	0.11	0.42	0.69	0.43
	$\widehat{\beta}_{0,i}^{FH+PS}$	0.27	-0.32	-0.03	0.26	0.56	0.87	0.48
Multi-strategy	Avg. $R_{i,t}$	0.48	0.11	0.27	0.44	0.67	0.96	0.37
	$\widehat{\beta}_{0,i}^{FF}$	0.26	-0.19	0.00	0.26	0.47	0.78	0.37
	$\widehat{\beta}_{0,i}^{FH}$	0.38	-0.10	0.20	0.37	0.57	0.80	0.36
	$\widehat{\beta}_{0,i}^{FH+PS}$	0.42	-0.08	0.18	0.40	0.62	0.90	0.39
Event Driven	Avg. $R_{i,t}$	0.48	0.00	0.29	0.49	0.70	0.83	0.36
	$\widehat{\beta}_{0,i}^{FF}$	0.21	-0.21	-0.01	0.20	0.44	0.60	0.37
	$\widehat{\beta}_{0,i}^{FH}$	0.34	-0.08	0.07	0.35	0.55	0.78	0.35
	$\widehat{\beta}_{0,i}^{FH+PS}$	0.45	-0.01	0.18	0.49	0.71	0.88	0.38
Relative Value	Avg. $R_{i,t}$	0.43	-0.02	0.24	0.47	0.63	0.82	0.35
	$\widehat{\beta}_{0,i}^{FF}$	0.40	-0.05	0.21	0.40	0.60	0.76	0.33
	$\widehat{\beta}_{0,i}^{FH}$	0.46	0.11	0.24	0.46	0.66	0.85	0.34
	$\widehat{\beta}_{0,i}^{FH+PS}$	0.53	0.06	0.30	0.53	0.80	0.98	0.41
Credit	Avg. $R_{i,t}$	0.44	0.03	0.30	0.43	0.60	0.78	0.41
	$\widehat{\beta}_{0,i}^{FF}$	0.39	0.00	0.21	0.41	0.61	0.85	0.38
	$\widehat{\beta}_{0,i}^{FH}$	0.47	0.07	0.30	0.49	0.72	0.90	0.36
	$\widehat{\beta}_{0,i}^{FH+PS}$	0.55	0.14	0.37	0.55	0.74	0.95	0.36
Macro	Avg. $R_{i,t}$	0.22	-0.24	0.10	0.29	0.43	0.54	0.40
	$\widehat{\beta}_{0,i}^{FF}$	0.02	-0.34	-0.15	0.10	0.25	0.39	0.36
	$\widehat{\beta}_{0,i}^{FH}$	0.10	-0.31	-0.09	0.15	0.38	0.48	0.36
	$\widehat{\beta}_{0,i}^{FH+PS}$	0.14	-0.28	-0.05	0.19	0.42	0.62	0.44
Other	Avg. $R_{i,t}$	0.49	-0.13	0.20	0.44	0.83	1.05	0.49
	$\widehat{\beta}_{0,i}^{FF}$	0.23	-0.22	0.00	0.21	0.51	0.82	0.46
	$\widehat{\beta}_{0,i}^{FH}$	0.27	-0.22	0.00	0.23	0.54	0.88	0.47
	$\widehat{\beta}_{0,i}^{FH+PS}$	0.32	-0.21	0.02	0.23	0.69	1.04	0.58

Table 4 reports summary statistics in total and by strategy for raw monthly returns and the intercepts in a first-stage time-series regression of returns on risk factors from three factor models: Fama-French-Carhart, Fung-Hsieh, and Fung-Hsieh plus Pastor-Stambaugh. The Managed Futures strategy are excluded due to insufficient observations.

Table 5: Parameter Values from Two-Stage Model without Illiquidity

		Percentiles						
FM Parameter Estimates	**Mean**	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}	**St. Dev**	**t-stat**
$\widehat{\phi}_{eq-mf}$	0.59	-2.35	-1.25	0.94	2.43	4.00	2.95	1.47
$\widehat{\phi}_{szspr-f}$	0.36	-3.39	-2.17	0.90	2.08	3.13	2.79	0.95
$\widehat{\phi}_{bd-mf}$	0.66	-21.01	-8.64	-0.73	11.84	24.62	21.02	0.23
$\widehat{\phi}_{crsprd-f}$	1.03	-16.60	-8.37	-0.50	12.83	26.63	17.87	0.42
$\widehat{\phi}_{agg-liq}$	-0.01	-0.06	-0.03	0.00	0.02	0.03	0.04	-1.84
$\widehat{\phi}_{innov-liq}$	-0.01	-0.06	-0.04	-0.02	0.03	0.05	0.05	-1.47
$\widehat{\phi}_{liq-v}$	0.00	-0.03	-0.01	0.00	0.02	0.03	0.03	0.00
$\widehat{\phi}_{bd-tff}$	-2.22	-21.24	-13.22	-2.98	8.19	15.39	15.14	-1.08
$\widehat{\phi}_{crrcy-tff}$	-0.73	-23.16	-14.24	-7.01	9.11	25.12	21.30	-0.25
$\widehat{\phi}_{cmdty-tff}$	-2.05	-23.05	-15.11	-2.57	7.58	17.54	17.90	-0.84

Table 5 reports the distribution of estimated parameters from equation (10).

Table 6: Correlation of Estimated Betas and Illiquidity

	$\overline{\log}(L_p)$	$\overline{\log}(L_s)$	$\widehat{\beta}_{bd\_tff}$	$\widehat{\beta}_{crrcy\_tff}$	$\widehat{\beta}_{cmdty\_tff}$	$\widehat{\beta}_{eq\_mf}$	$\widehat{\beta}_{szspr\_f}$	$\widehat{\beta}_{bd\_mf}$	$\widehat{\beta}_{crsprd\_f}$	$\widehat{\beta}_{agg\_liq}$	$\widehat{\beta}_{innov\_liq}$	$\widehat{\beta}_{liq\_v}$
$\overline{\log}(L_p)$	1.00											
$\overline{\log}(L_s)$	0.51	1.00										
$\widehat{\beta}_{bd\_tff}$	-0.07	-0.05	1.00									
$\widehat{\beta}_{crrcy\_tff}$	-0.06	-0.02	-0.50	1.00								
$\widehat{\beta}_{cmdty\_tff}$	-0.03	-0.05	0.07	-0.29	1.00							
$\widehat{\beta}_{eq\_mf}$	-0.21	-0.12	-0.12	0.06	0.01	1.00						
$\widehat{\beta}_{szspr\_f}$	0.14	0.19	-0.13	-0.13	0.08	0.26	1.00					
$\widehat{\beta}_{bd\_mf}$	0.24	0.33	-0.08	0.04	-0.18	0.01	0.10	1.00				
$\widehat{\beta}_{crsprd\_f}$	0.01	0.16	-0.36	0.21	-0.05	0.14	0.09	0.53	1.00			
$\widehat{\beta}_{agg\_liq}$	0.10	0.12	-0.19	0.15	-0.25	0.26	0.30	0.15	0.03	1.00		
$\widehat{\beta}_{innov\_liq}$	-0.08	-0.10	0.13	-0.10	0.12	-0.25	-0.18	-0.16	-0.09	-0.80	1.00	
$\widehat{\beta}_{liq\_v}$	0.04	-0.01	0.24	-0.12	-0.17	0.11	-0.24	0.08	0.02	0.00	-0.04	1.00

Table 6 reports the unconditional correlations between the betas estimated in the first-stage of the Fama-MacBeth procedure (equation 9), and the measures of portfolio and investor share illiquidity, $\overline{\log}(L_p)$ and $\overline{\log}(L_s)$. The table includes only one observation per fund.

Table 7: Distribution of Alphas and Mean Absolute Pricing Errors

		Percentiles					
Fama-French-Carhart	**Mean**	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}	**St. Dev**
α_i^{FF}	0.33	-0.10	0.11	0.33	0.55	0.76	0.38
α_i^{s}	0.03	-0.44	-0.19	0.05	0.26	0.47	0.38
α_i^{p}	0.17	-0.29	-0.05	0.17	0.39	0.60	0.38
$\lvert\alpha_i^{FF}\rvert$	0.41	0.07	0.19	0.37	0.58	0.78	0.29
$\lvert\alpha_i^{s}\rvert$	0.29	0.04	0.11	0.23	0.41	0.63	0.25
$\lvert\alpha_i^{p}\rvert$	0.32	0.05	0.13	0.27	0.44	0.66	0.26
Fung-Hsieh							
α_i^{FH}	0.37	-0.06	0.17	0.37	0.58	0.80	0.37
α_i^{s}	0.13	-0.31	-0.05	0.14	0.34	0.55	0.37
α_i^{p}	0.23	-0.22	0.04	0.23	0.43	0.66	0.37
$\lvert\alpha_i^{FH}\rvert$	0.44	0.10	0.23	0.40	0.59	0.81	0.29
$\lvert\alpha_i^{s}\rvert$	0.30	0.04	0.10	0.24	0.42	0.61	0.26
$\lvert\alpha_i^{p}\rvert$	0.34	0.06	0.15	0.29	0.48	0.69	0.27
Fung-Hsieh + Pastor-Stambaugh							
α_i^{FH+PS}	0.39	-0.06	0.17	0.38	0.59	0.81	0.37
α_i^{s}	0.14	-0.30	-0.05	0.15	0.35	0.54	0.37
α_i^{p}	0.24	-0.20	0.05	0.25	0.45	0.67	0.37
$\lvert\alpha_i^{FH+PS}\rvert$	0.45	0.10	0.23	0.41	0.60	0.82	0.30
$\lvert\alpha_i^{s}\rvert$	0.30	0.04	0.11	0.24	0.42	0.63	0.25
$\lvert\alpha_i^{p}\rvert$	0.35	0.06	0.15	0.30	0.49	0.68	0.27
Fung-Hsieh + Pastor-Stambaugh GLM Liquidity (θ_0)							
$\alpha_i^{\theta_0}$	0.33	-0.09	0.15	0.32	0.53	0.75	0.36
$\lvert\alpha_i^{\theta_0}\rvert$	0.41	0.10	0.21	0.36	0.55	0.77	0.27

Table 7 reports the distribution of estimated alphas and mean absolute pricing errors from equation 14. Only one observation is included per fund.

Table 8: FH+PS Alphas: Gross vs. Net Returns

Fung-Hsieh + Pastor-Stambaugh	Mean	Percentiles 10th	25th	50th	75th	90th	St. Dev
α_i^{FH+PS} (Net)	0.39	-0.06	0.17	0.38	0.59	0.81	0.37
α_i^{FH+PS} (Gross)	0.56	0.07	0.30	0.56	0.80	1.07	0.44
α_i^{s} (Net)	0.14	-0.30	-0.05	0.15	0.35	0.54	0.37
α_i^{s} (Gross)	0.25	-0.26	0.00	0.24	0.49	0.74	0.43
α_i^{p} (Net)	0.24	-0.20	0.05	0.25	0.45	0.67	0.37
α_i^{p} (Gross)	0.41	-0.10	0.17	0.39	0.65	0.91	0.43

Table 8 reports the distribution of estimated alphas and mean absolute pricing errors from equation 14. Only one observation is included per fund.

Table 9: Share-Portfolio Mismatch vs. Financing Liquidity

	Share-Portfolio Mismatch	Share-Portfolio Mismatch	$\overline{L}_p$	Share-Portfolio Mismatch	Share-Portfolio Mismatch	Share-Portfolio Mismatch
$\overline{L}_f$	0.090**	0.103**				0.087**
	(0.041)	(0.042)				(0.041)
Event Driven (Equity)			-25.595	65.311**	63.833**	63.216**
			(22.029)	(27.752)	(27.769)	(27.729)
Constant	106.485***	101.186***	42.462***	109.020***	107.654***	104.481***
	(3.278)	(11.233)	(2.326)	(2.930)	(4.157)	(4.414)
Strategy Dummies	No	Yes	No	No	No	No
Leverage Controls	No	Yes	No	No	Yes	Yes
Size Controls	No	Yes	No	No	Yes	Yes
Adjusted R^2	0.003	0.052	0.000	0.004	0.004	0.007
Observations	1166	1166	1166	1166	1166	1166

Table 9 reports regression results of the share-portfolio mismatch, defined as $\overline{L}_s - \overline{L}_p$, on average financing illiquidity $\overline{L}_f$ and a dummy variable for whether the fund characterizes its strategy as event-driven. All regressions in this table are cross-sectional. Size controls include net asset value, leverage controls include the ratios of gross assets, borrowing, and gross notional exposures to net assets, and strategy dummies are based on broad strategy category. Regressions are run only for the 1,166 funds with estimated alpha values and non-zero borrowing. * denotes significance at the 10% level, ** at the 5% level, and *** at the 1% level.

Table 10: FH+PS Alphas vs. Autocorrelation Coefficients for 5 Illiquidity Portfolios

Sorted Portfolio	Avg. $\overline{L}_s$	ρ_p	ρ_f	$\overline{R}_i$	$\widehat{\alpha}_{FH+PS}$
1 (Liquid)	11	0.012	0.012	0.421	0.207
2	57	0.043	0.060	0.530	0.336
3	122	0.106	0.121	0.506	0.374
4	231	0.275	0.110	0.515	0.384
5 (Illiquid)	344	0.293	0.136	0.647	0.588

Table 10 reports average statistics for funds divided into five groups based on their average investor share illiquidity. Statistics shown are the investor share illiquidity, the first-order autocorrelation coefficients of fund returns, computed based on returns averaged within each group (ρ_p) and based on the average of the fund specific return autocorrelation coefficients (ρ_f), the average raw monthly returns within each group, and the average FH+PS alphas within each group.

Table 11: Alphas, Liquidity, and Fees

	Avg. Fee	Avg. Fee	Avg. Fee	Avg. Fee	Avg. Fee	Incentive Fee	Incentive Fee	Management Fee
$a_i^{s,g}$	1.333***	1.266***	0.755***	0.954***	0.937***		3.120***	0.300***
	(0.112)	(0.111)	(0.123)	(0.123)	(0.122)		(0.747)	(0.098)
$\overline{\log}(L_p)$		0.159***	0.112***	0.014	0.026			
		(0.027)	(0.027)	(0.029)	(0.029)			
$\% \ r_{i,t} > 0$			2.961***	2.201***	2.132***			
			(0.343)	(0.350)	(0.347)			
$\overline{\log}(L_s)$				0.232***	0.229***		1.977***	0.034
				(0.031)	(0.031)		(0.213)	(0.026)
Avg. Leverage					0.067***		0.323***	0.025
					(0.016)		(0.105)	(0.016)
Man. Fee						2.054***		
						(0.316)		
Constant	2.215***	1.711***	-0.264	-0.583**	-0.659**	12.758***	3.512**	1.273***
	(0.163)	(0.182)	(0.289)	(0.285)	(0.283)	(0.471)	(1.596)	(0.208)
Size Controls	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
Strategy Controls	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
Adjusted R^2	0.177	0.204	0.257	0.294	0.306	0.055	0.169	0.015
Observations	1037	1037	1037	1037	1037	710	710	1055

Table 11 shows regressions of the average realized fee, estimated incentive fee, and estimated management fee on estimated alpha, investor share illiquidity, and other controls. $\#r_{i,t} > 0$ is the number of months in the sample for which the fund had a positive gross rate of return. Size controls include net asset value, leverage controls include the ratios of gross assets, borrowing, and gross notional exposures to net assets, and strategy dummies are based on broad strategy category. * denotes significance at the 10% level, ** at the 5% level, and *** at the 1% level.

Table 12: Alphas, Liquidity, and Fees

	Avg. Fee	Avg. Fee	Avg. Fee	Avg. Fee	Avg. Fee	Incentive Fee	Incentive Fee	Management Fee
a_i^s	0.753***	0.666***	-0.111	0.062	0.064		0.449	-0.056
	(0.135)	(0.133)	(0.141)	(0.142)	(0.141)		(0.888)	(0.115)
$\overline{\log}(L_p)$		0.176***	0.110***	0.034	0.046			
		(0.028)	(0.027)	(0.030)	(0.030)			
$\% \ r_{i,t} > 0$			4.109***	3.529***	3.426***			
			(0.345)	(0.356)	(0.353)			
$\overline{\log}(L_s)$				0.182***	0.180***		1.973***	0.033
				(0.032)	(0.032)		(0.216)	(0.026)
Avg. Leverage					0.071***		0.335***	0.028*
					(0.016)		(0.106)	(0.016)
Man. Fee						2.054***		
						(0.316)		
Constant	2.508***	1.942***	-0.876***	-1.118***	-1.187***	12.758***	4.638***	1.386***
	(0.169)	(0.190)	(0.296)	(0.295)	(0.293)	(0.471)	(1.601)	(0.208)
Size Controls	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
Strategy Controls	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
Adjusted R^2	0.092	0.124	0.230	0.252	0.266	0.055	0.149	0.006
Observations	1037	1037	1037	1037	1037	710	710	1055

Table 12 shows regressions of the average realized fee, estimated incentive fee, and estimated management fee on estimated alpha, investor share illiquidity, and other controls. $\#r_{i,t} > 0$ is the number of months in the sample for which the fund had a positive gross rate of return. Size controls include net asset value, leverage controls include the ratios of gross assets, borrowing, and gross notional exposures to net assets, and strategy dummies are based on broad strategy category. * denotes significance at the 10% level, ** at the 5% level, and *** at the 1% level.

8.2 Figures

Figure 1: Distribution of L_p

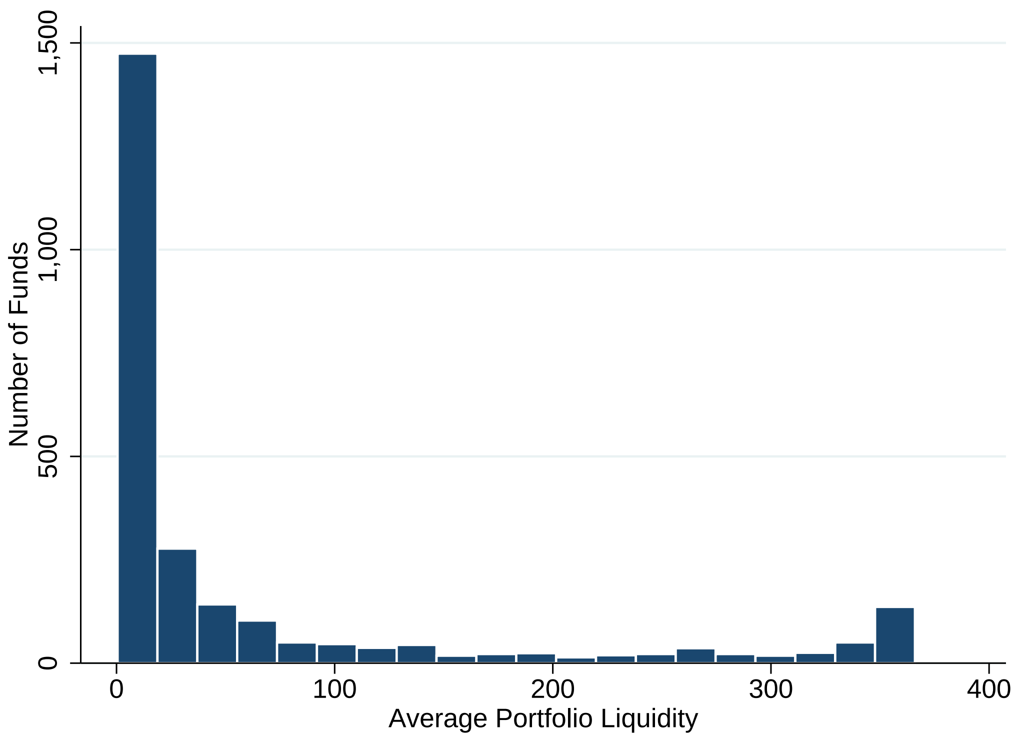

Figure 1 shows a histogram plot of the distribution of average portfolio illiquidity, $\overline{L}_p$, for our main sample of 2,518 funds.

Figure 2: Distribution of L_s

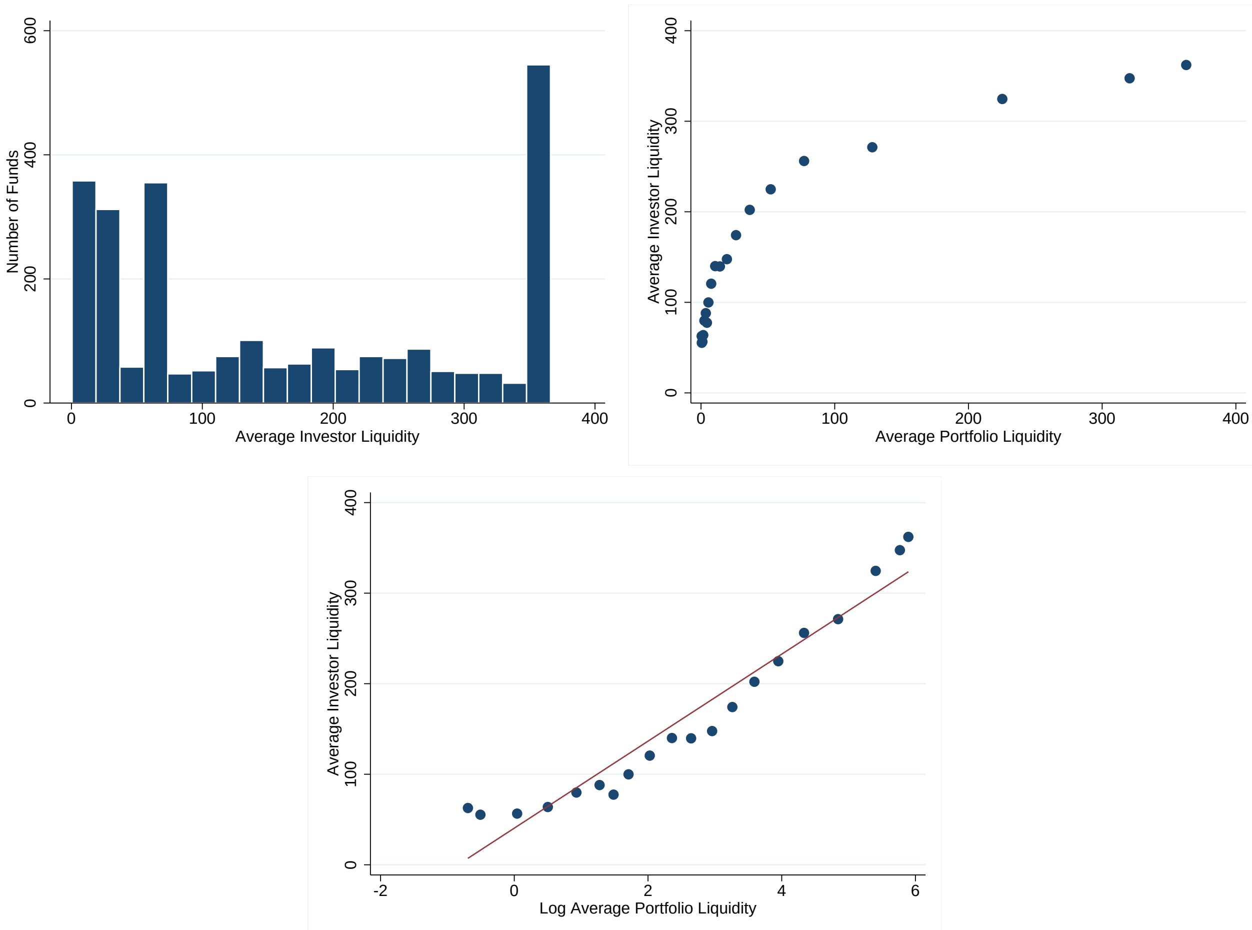

Figure 2 shows a average investor share illiquidity, $\overline{L}_s$, along with average portfolio illiquidity, $\overline{L}_p$, for our main sample of 2,518 funds. The upper left panel plots the histogram of $\overline{L}_s$. The upper right panel shows a bin-scatter plot of $\overline{L}_s$ versus $\overline{L}_p$, and the bottom panel shows a bin-scatter plot of $\overline{L}_s$ versus $\log(\overline{L}_p)$.

Figure 3: R^2 in Fama-MacBeth Second Stage Regressions (FH+PS factors)

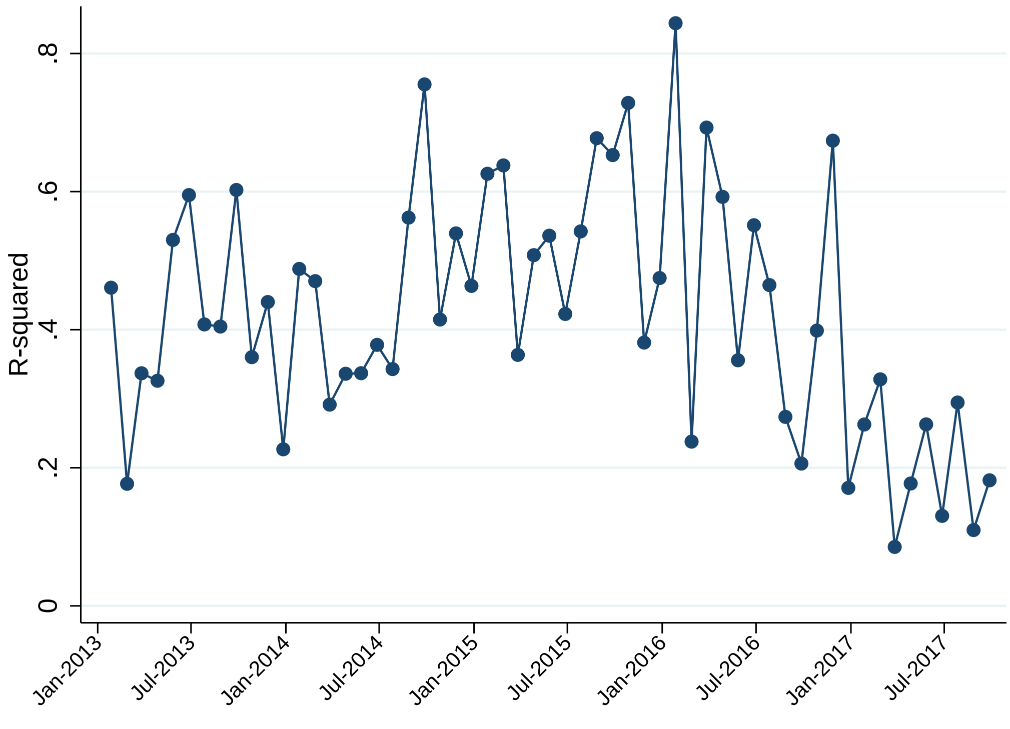

Figure 3 plots the R-squared from each of the T second-stage regressions estimated in equation (10).

Figure 4: Alphas Estimated From Time-Series and Two-Pass Regressions

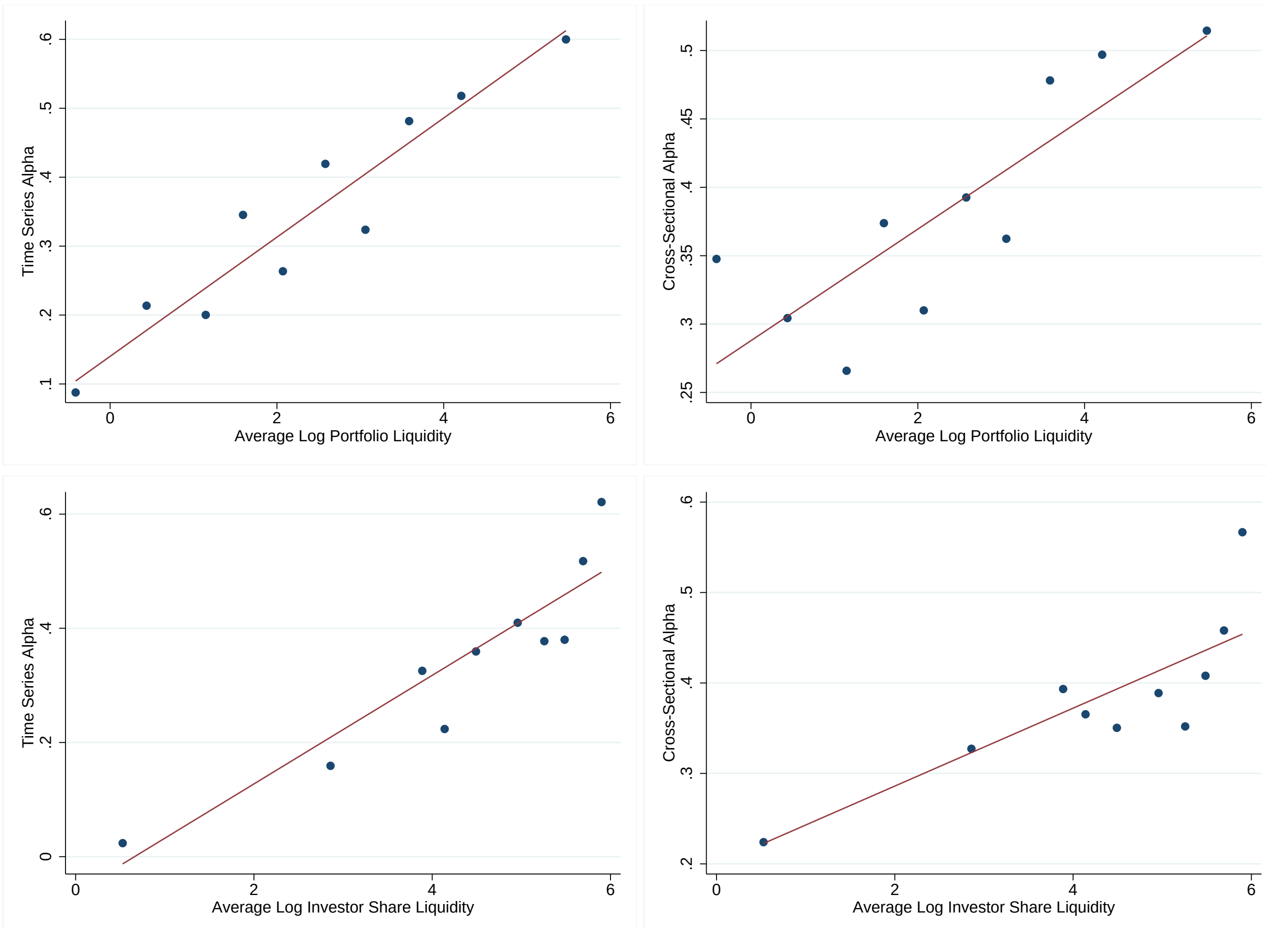

Figure 4 shows bin-scatter plots of estimated alphas versus $\overline{\log}(L_p)$ and $\overline{\log}(L_s)$. Time series alpha corresponds to β_0^{FH+PS} estimated in equation (9), and cross-sectional alpha refers to α_i^{FH+PS} estimated in equation (14). Funds are separated into 10 groups based on either $\overline{\log}(L_p)$ or $\overline{\log}(L_s)$, and means for estimated alphas and each measure of illiquidity are plotted.

Figure 5: Illiquidity Premium

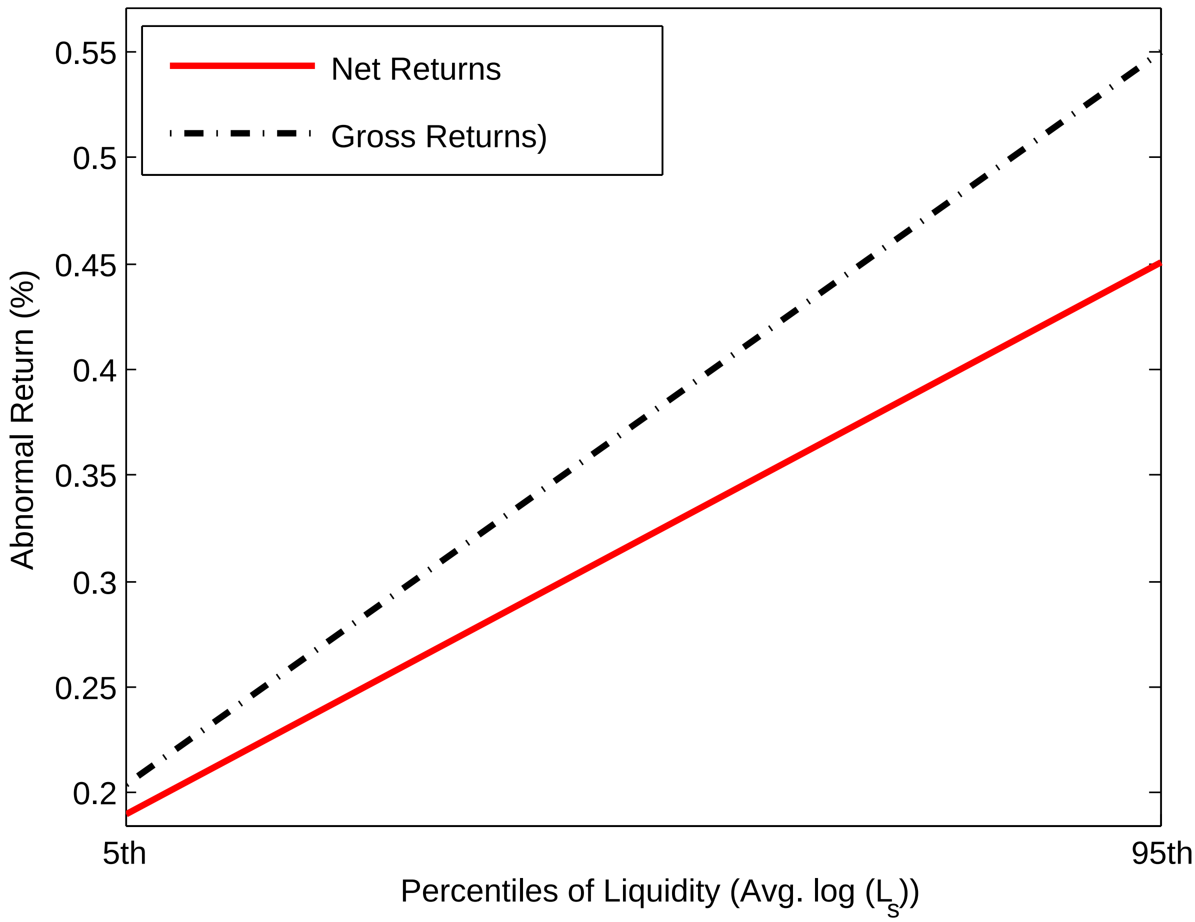

Figure 5 shows the estimated illiquidity premium $\widehat{\lambda}$ from equation (14), using both gross-of-fee and net-of-fee returns on the left-hand side. The horizontal axis is investor share illiquidity, represented by the corresponding percentiles. The vertical axis is the implied *ceteris paribus* excess monthly return associated with investor share illiquidity; it is based on the median value of α^{FH+PS} and $\widehat{\lambda}$.

Figure 6: Illiquidity Risk Factor

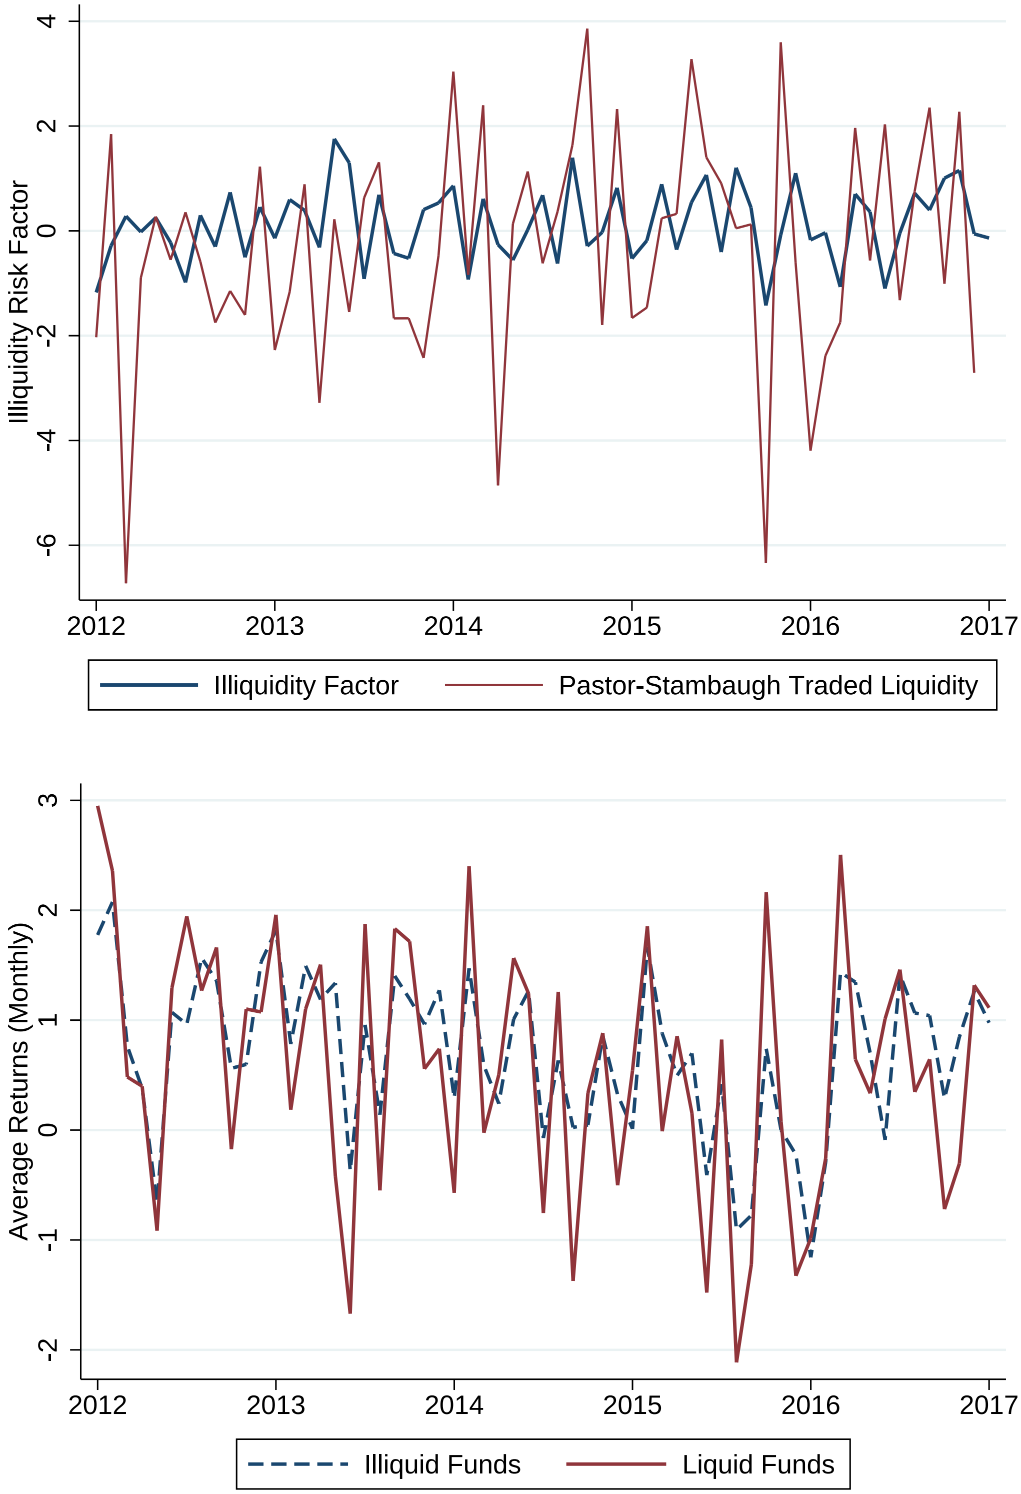

The top panel of Figure 6 shows the time-series of the illiquidity risk factor versus the traded liquidity factor from Pastor and Stambaugh (2003) (normalized to be on equivalent scales). The bottom panel shows the time series of average monthly returns for the 20% most and least liquid funds in Form PF.

Figure 7: Estimated Fees

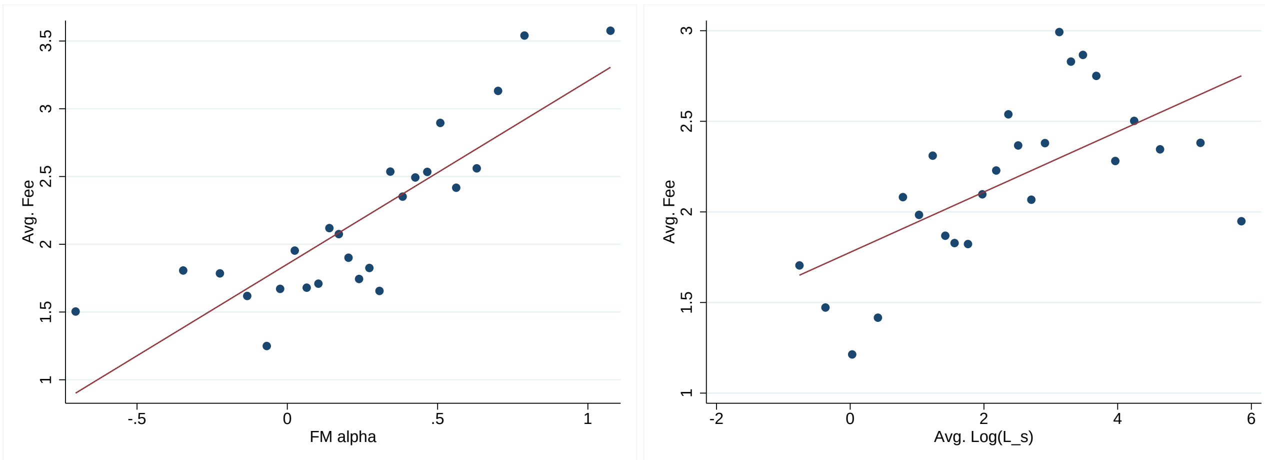

Figure 7 shows bin-scatter plots of average realized fees, calculated as the reported gross rate of return minus the net rate of return. The left-panel plots realized fees against values of α_i^s (FM alpha). The right panel plots realized fees against $\overline{\log}(L_s)$, controlling for α_i^s, so that the values of realized fees on the vertical axis are residualized with respect to α_i^s.