OFFICE OF FINANCIAL RESEARCH
WORKING PAPER

20-03 | February 25, 2020

Illiquidity in Intermediate Portfolios: Evidence from Large Hedge Funds

Daniel Barth
Office of Financial Research
daniel.barth@ofr.treasury.gov

Phillip Monin
Office of Financial Research
phillip.monin@ofr.treasury.gov

Illiquidity in Intermediary Portfolios: Evidence from Large Hedge Funds

Daniel Barth,[1,2,3] Phillip Monin[4]

October 22, 2019

Abstract

Evaporating liquidity is a central feature of many financial crises. Questions remain about the importance of illiquidity and the distribution of illiquidity exposure across financial market participants. We use regulatory data on hedge funds — who unlike public mutual funds often invest in illiquid markets — to address three empirical questions: how large is the illiquidity premium, how important is it for hedge fund returns, and who ultimately captures the premium, fund managers or investors? We estimate an annual illiquidity premium of 56 basis points for an additional log-day needed to sell assets without price impact, of which investors capture 77%. Portfolio illiquidity explains 27% of alpha, but share restrictions explain 55%. Consistent with compensation for undiversifiable illiquidity risk, managers of illiquid funds charge higher incentive fees. Our findings suggest the costs and risks associated with illiquid assets are substantial and require significant compensation. Moreover, the returns of some funds are highly dependent on the illiquidity premium, which may indicate these funds have significant exposures to illiquidity. However, through share restrictions much of this exposure is passed to fund investors, who are likely better able to diversify across many asset classes.

JEL Classifications: G11, G12, G23

Keywords: Hedge Funds, Illiquidity, Share Restrictions, Intermediaries

---

[1] 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.

[2] We are thankful to George Aragon, Zahi Ben-David, Greg Brown, Greg Duffee, Bing Liang, Christian Lundblad, Mila Getmansky Sherman, Zhen Shi (WFA discussant), Stephen Shore, Emil Siriwardane, Russ Wermers, and seminar participants at the 2018 Western Finance Association Annual Meetings, 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 for valuable comments and discussions. All remaining errors are our own.

[3] 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

[4] 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

Widespread illiquidity is a hallmark of financial crises (Allen and Gale (2000), Bernardo and Welch (2004)). Illiquid assets become difficult to sell during times of market stress, forcing quick and possibly severe devaluations and potentially instigating liquidity spirals that destabilize financial markets (Brunnermeier and Pedersen (2009)). However, nearly all measures of illiquidity, such as price impact per volume or short-term reversals, are based on data from trades.[5] That is, existing estimates of illiquidity and the illiquidity premium are determined only from the most liquid assets.

Yet, the most relevant implications for systemic risk likely reside in deeply illiquid markets, those with assets that trade infrequently or in some cases not at all. In this paper, we pursue an alternative methodology to study illiquidity based on the portfolios of traders who are likely to be meaningful participants in illiquid markets: hedge funds.[6] The returns to hedge funds' portfolios, which often comprise investments in illiquid assets, offer a way to measure illiquidity and its premium from the full spectrum of illiquidity, including deeply illiquid assets for which trade data may not be available.

Specifically, we address three empirical questions related to illiquidity in hedge funds: how large is the illiquidity premium, how important is it for explaining hedge funds' risk-adjusted returns, and who ultimately captures the premium, the fund manager or fund investors?[7] The size of the illiquidity premium is relevant to financial stability because it represents market participants' view of the importance of illiquidity; that is, it provides information on the magnitude of the costs and risks associated with holding illiquid assets. The dependence of hedge funds' returns on illiquidity matters because it suggests the extent to which funds are exposed to illiquidity risk. If funds' returns are not reliant on illiquidity, then funds must have material illiquidity exposure. Finally, who captures the premium matters because those who bear the costs and risks of illiquid assets are the ones who earn the premium, and fund investors are likely better able to diversify exposure to illiquidity than fund managers because managers often have a large personal stake in

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

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

[7]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.

the fund.

We find that the illiquidity premium is large: an additional log-day needed to sell assets without price impact is associated with an extra 0.56% per year in risk-adjusted returns. The illiquidity premium is also important for explaining hedge fund returns: portfolio illiquidity explains 27% of average estimated alpha, whereas investor share illiquidity explains 55%. Finally, we find that investors capture most, but not all, of the illiquidity premium. Investors receive 77% of the premium associated with illiquid assets, and receive a similar amount of the premium associated with the illiquidity of their shares. We offer an economic interpretation for these findings based on the premise that the level of illiquidity of fund assets constitutes an undiversifiable risk for the fund manager.

Importantly, our measure of portfolio illiquidity is based on expected liquidation horizons reported by fund managers themselves, rather than on proxies such as the autocorrelation of returns. Our data comprise a rich panel of hedge fund activities that originates from the first systematic, regulatory collection on large private 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, investor share restrictions, the term structure of financing, and much more. Form PF offers an unprecedented view of the activities of large hedge funds, which often do not report to any of the public databases.

Our empirical analyses are motivated by an economic argument for who ultimately bears the costs of holding illiquid assets, the hedge fund manager or hedge fund investors, and therefore who ultimately captures the illiquidity premium. The unifying theory that gives economic content to our findings is that the level of illiquidity of fund assets constitutes an undervisifiable risk for the hedge fund manager. This arises because (i) due to principle-agent concerns, investors require managers to invest a substantial portion of their personal wealth in the fund, and (ii) the forced sale of illiquid assets, either due to investor redemptions or margin calls on leveraged positions, is likely to result in price impact. The cost of this price impact is borne by the (remaining) fund investors, of whom the manager is the ultimate residual claimant. The manager is unable to diversify this risk because of his large personal investment in the fund's assets.

If fund managers are unable to fully diversify the possible price impact associated with investor redemptions, they can respond either by requiring explicit compensation for this risk in the form of higher fees, or by passing the costs to investors by restricting investor shares (or both). This insight generates a number of empirical predictions. If managers retain the entirety of the illiquidity risk — implying they offer on-demand redemptions and are compensated with higher fees — the illiquidity premium will be found in gross-of-fee but not net-of-fee returns. In this case, investors bear none of the costs of illiquid assets, and therefore capture none of the premium. Further, funds' asset illiquidity will be positively correlated with fees but will be uncorrelated 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 (and the risk of forced sales from margin calls is negligible), 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. If managers and investors share the illiquidity premium, or if margin risk is not negligible, the illiquidity premium should appear in both net and gross returns. In this case, the difference in the illiquidity premium between gross and net returns will determine the proportion of the illiquidity premium passed to fund investors. Moreover, both fees and share restrictions will be positively correlated with asset illiquidity. To our knowledge, this hypothesized relationship between asset illiquidity, share restrictions, and fees has not been proposed elsewhere in the literature.

This argument is rooted in theories of limits to arbitrage, in which the marginal investors in certain asset markets are undiversified specialists rather than well-diversified traders. Hedge funds are often considered the quintessential examples of such specialists. These theories predict that the liquidation of assets in these markets can have a large effect on prices, due to the smaller and more specialized pool of investors (Shleifer and Vishny (1997), Gromb and Vayanos (2002), Kyle and Xiong (2001), He and Krishnamurthy (2012), Gabaix, Krishnamurthy, and Vigneron (2007), Siriwardane (2019)). These theories are supported by empirical evidence of a large and persistent price impact from forced sales by institutional traders (Coval and Stafford (2007), Jotikasthira, Lundblad, and Ramadorai (2012), Falato, Hortascu, Li, and Shin (2017)).

Motivated by our economic argument, our first set of results establishes an 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 needed to meet investor redemptions. 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 approximately 0.50 and has a robust t-statistic of 34, which suggests that share illiquidity increases more quickly than asset illiquidity on average. This offers evidence that at least some of the costs associated with illiquid assets ultimately reside with the hedge fund investor rather than the fund manager.

Next, we measure the illiquidity premium implied by hedge fund portfolios. The nature of the data allow for an economically intuitive interpretation of the premium — the additional return due to an extra log-day needed for managers to liquidate assets without price impact or for investors to redeem their shares. Our estimates imply that an additional log-day of portfolio illiquidity is associated with an excess return of 56 basis points per year. Alternatively, the *ceteris paribus* excess return earned by a fund in the 95$^{\text{th}}$ percentile of portfolio illiquidity is 3.4% (percentage points) higher per year than a fund in the 5$^{\text{th}}$ percentile. The magnitudes of these premiums suggest that the costs and risks of illiquid assets are substantial, and require commensurate compensation from financial market participants.

Similarly, an additional log-day of investor share illiquidity produces an excess return of 76 basis points per year. This suggests share restrictions may provide economic value to investors beyond matching the illiquidity of assets and liabilities, such as greater managerial discretion (Agarwal, Daniel, and Naik (2009)), an ability to exploit opportunistic trading opportunities during times of market stress and tight funding conditions (Shleifer and Vishny (1997), Aragon, Martin, and Shi (2019)), or to pursue profitable trades with uncertain payoff horizons (Giannetti and Kahraman (2017)). Or, it may be that the measure of average portfolio illiquidity we use fails to account for the entirety of illiquidity exposure, and tighter share restrictions better proxy for such exposure.

We then estimate the importance of illiquidity for explaining hedge funds' risk-adjusted returns. In this context our data are particularly valuable because our measures of illiquidity span much longer horizons than existing measures inferred from trading data. Based on standard Fama-MacBeth regressions, we find that average alphas estimated from traditional factor models and based on gross-of-fee returns are 27–39% smaller after controlling for portfolio illiquidity. That is, exposure to illiquidity comprises a significant portion of hedge fund returns, which may indicate hedge funds are highly exposed to the risks associated with illiquidity assets.

A similar analysis finds that investor share illiquidity has considerably more explanatory power, reducing average estimated alphas by over 55-80%, depending on the model. We also find portfolio and investor share illiquidity reduce mean average pricing errors by between 19% and 38%, again depending on the model. While gross-of-fee returns are higher, we find reductions of alphas and pricing errors of similar magnitudes when net-of-fee returns are used instead. We explore whether return smoothing, a well-established characteristic of hedge fund returns (Asness, Krail, and Liew (2001), Khandani and Lo (2011)) explains the size of the estimated illiquidity premium or reductions in estimated alphas, but find no evidence that it has a meaningful effect on our results.

Based on the differences in risk-adjusted returns between the most and least liquid funds, we next examine whether our measure of portfolio illiquidity — which is properly interpreted as a fund characteristic — may also represent an asset pricing risk factor. We find no evidence in support of this hypothesis. We construct an illiquidity factor based on mean return differences in each period between the 20% most illiquid and liquid funds. While this illiquidity risk factor has a 32% correlation with the Pastor and Stambaugh (2003) traded liquidity factor, suggesting it has some information content, the estimated price of risk is economically small and statistically insignificant. This suggests that while the illiquidity premium may in part represent compensation for fund managers who cannot effectively diversify price impact risks, it is not compensation for exposure to a systematic risk factor.

Finally, we investigate who ultimately captures the illiquidity premium. The return premiums estimated from gross and net returns separately also allow us to estimate the average fraction of the illiquidity premium passed through to fund investors. We estimate that about 77% of the portfolio illiquidity premium earned by hedge funds is passed through to fund investors, and estimate a slightly smaller pass-through rate of 72% for the investor share illiquidity premium. We find a similar pass-through rate on the investor share illiquidity premium when both portfolio and investor share illiquidity are included in the regression simultaneously. The finding that most of the illiquidity risk is passed to fund investors may have important implications for financial stability; investors are likely better able to diversify across many strategies and asset classes than fund managers, and may therefore be less likely to demand liquidity in the event of market stress.

The pass-through rate implied by gross and net returns is informative, but also problematic. If all funds charged “2 and 20”, then the illiquidity premium pass-through rate would be 20% by construction. Our economic framework suggests specifically that the *rate* of fees should be associated with illiquidity.[8] Further, theories of intermediary-based asset pricing (He and Krishnamurthy (2012), He and Krishnamurthy (2013)) predict the relationship between fees and illiquidity should emerge only through the incentive fee.

Consistent with managers retaining some of the risk associated with illiquid assets and requiring explicit compensation for this risk, we show that realized fees are positively associated with portfolio illiquidity. These results are robust to a host of controls intended to proxy for alternative hypotheses. Further, as predicted by our economic argument, we also show the relationship between realized fees and illiquidity is entirely driven by the incentive fee, which scales up proportionately with the illiquidity premium. Our last result is perhaps the strongest evidence in support of our economic argument; we find that among funds that are 0% manager-owned, the relationship between illiquidity and fees completely disappears. That is, in funds where managers have no personal wealth stake and therefore no undiversifiable exposure to illiquidity, liquid and illiquid funds charge similar fees.

Two existing papers also study liquidity using Form PF, each of which uses a similar summary measure of portfolio and investor share liquidity as the ones used here (Aragon, Ergun, Getmansky, and Girardi (2017b) and Aragon, Ergun, Getmansky, and Girardi (2017a)). While those papers focus on the liquidity characteristics and risk management practices of hedge funds, including the analysis of liquidity buffers and the liquidity mismatch between assets and liabilities, our paper focuses on the size and importance of the illiquidity premium. Measures of liquidity mismatch and the illiquidity premium are each useful for a more complete understanding of systemic risk; liquidity mismatch determines the amount of available liquidity (or lack thereof) in cases of unanticipated liquidity needs, such as redemptions or margin calls, and the illiquidity premium offers a way to measure the size and importance of illiquidity exposure in hedge funds and financial markets more generally.

[8] Our analysis of hedge fund fees is not concerned with their absolute levels or trends, or their appropriateness in contracting between fund managers and investors. Rather, we are interested in fees solely with regards to their relationship with illiquidity exposure.

Our paper is also 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 several important ways. First, the TASS database primarily comprises smaller funds; our data on large funds is much more representative of the industry based on total assets under management. Second, Aragon (2007) uses investor share restrictions as a proxy for portfolio illiquidity. However, we are able to estimate illiquidity premia from portfolio and investor share illiquidity separately. Further, using data on gross and net returns, we are able to determine how much of the illiquidity premium is ultimately captured by fund 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).[9] They estimate a premium of 3.96% across all funds, and much larger premiums within strategy. We find, however, that return smoothing is only partially correlated with illiquidity in our sample, and that return premiums estimated from autocorrelations in our Fama-MacBeth setup 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 relationships among hedge fund returns, portfolio illiquidity, investor share illiquidity, and fees. Section 3 describes our data and provides fund-level summary statistics. Section

[9] 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.

4 provides hedge fund illiquidity summary statistics and establishes the empirical relationships between portfolio illiquidity and investor share illiquidity. Section 5 provides estimates of the illiquidity return premium and evaluates the importance of illiquidity in explaining hedge fund alphas. Section 6 analyzes the fraction of the premium captured by fund investors. Section 7 concludes.

2 The Economics of Illiquidity, Share Restrictions, and Fees

The empirical analyses performed in Sections 4–6 relate portfolio illiquidity, investor share restrictions, returns, and fees. In this section, we give these relationships an economic interpretation. The central premise is that the level of portfolio illiquidity constitutes an undiversifiable risk for the hedge fund manager. This premise, in concert with hedge funds' roles as intermediaries, produces a set of empirical predictions that we explore throughout the remainder of the paper.

2.1 The Risks of Illiquid Assets in Hedge Fund Portfolios

Modern asset pricing theory asserts that expected returns on risky securities are determined by their sensitivities to undiversifiable risk factors. Valid risk factors are constructed as returns to traded portfolios that represent a systematic risk. However, recent research on 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 (2019)). 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 assets that are difficult to price, the costs of information acquisition and the necessity of asset-specific expertise give rise to natural barriers to entry. This suggests that 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 them 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).

An important feature of hedge funds is that 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 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 price impact risk that arises in illiquid markets constitutes an undiversifiable risk for the fund manager.[10]

2.2 Empirical Predictions

For managed funds, immediate liquidity needs originate from two primary sources: financing liquidity shocks associated with margin calls or a sudden inability to rollover short-term debt, or 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 is described by a standard asset pricing model of the form estimated in Section 5:

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

[10]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.

That is, the expected return to the portfolio of fund j is determined by sensitivities to standard risk factors (β_j) multiplied by their respective prices of risk (γ), the illiquidity of fund assets L_j multiplied by 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 illiquidity, which are based on liquidation horizons of assets and redemption horizons of investors.

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, $\pi_{j,1}$. 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 an initial investor redemption request and when the fund actually meets that request, and a redemption frequency that establishes how often funds will pay redemptions (for example, quarterly). For now, we ignore margin risk and assume that the only funding liquidity risk managers face is from investor redemptions.

In this case, competition will force the entirety of the illiquidity premium (λ) to be passed to fund investors, because investors bear all the costs of illiquid assets:

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

This also implies that the incentive fee $\pi_{j,1}$ will be a function of L_j. For simplicity, assume portfolio illiquidity is the only source of fee heterogeneity: $\pi_{1,j} = \pi_1(L_j)$ and $\pi_{0,j} = \pi_0$. Then equation (3) implies:

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

If $E[R_j^{gross}] > 0$, $\lambda > 0$, and $\pi_1 > 0$, meaning that expected gross returns, the illiquidity premium, and the incentive fee 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 the incentive fee 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 do the share restrictions; this is a necessary condition for managers 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, then: (1) asset illiquidity and investor share restrictions will be tightly (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, the incentive fee and illiquidity must be positively correlated. This offers an alternative empirical hypothesis:

> *If fund managers pass none of the illiquidity risk associated with investor redemptions to fund investors, then: (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, then: (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. In this case the hedge fund investor becomes the marginal investor in illiquid asset markets rather than the manager, and the underlying illiquid assets will still command a return premium. This argument suggests 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 capacities 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.

3 Data

We use fund-level data from the Security and Exchange Commission's (SEC) 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.[11] Private fund advisers file at least annually and report items such as gross and net asset values, gross- and net-of-fee 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*.[12] 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

[11]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.

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

quarter.[13] In this paper, we focus exclusively on Qualifying Hedge Funds.[14]

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 that occur after 2012. 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, borrowing, net asset value (NAV), gross asset value (GAV), and fund strategy. We further require gross assets to be greater than or equal to net assets, and both to be greater than zero. We also exclude funds that primarily invest in other private funds, as well as funds in the bottom 1% of average NAV.

Whereas 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 main sample is an unbalanced panel of 2,550 funds for the period January 2013 to September 2017.[15] In total, we have 91,188 fund-month observations and 27,132 fund-quarter observations for funds in our main sample. When we investigate fund returns in Sections 5 and 6, we further restrict the sample to funds with first-stage alpha values (β_0) between the 2.5$^{\text{th}}$ and 97.5$^{\text{th}}$ percentiles for each of the three factor models.

The top panel of Table 1 reports cross-sectional summary statistics for our main sample. Funds have an average net asset value of \$1.14 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 \$2.06 billion but has a median of \$0.69 billion. Balance sheet leverage, defined as the ratio of gross to net assets, is on average 3.46 but has a median and 75th percentile of 1.17 and 1.68, respectively. Over the 19 quarters in our sample, we have a median of 10 quarterly

[13] 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.

[14] 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.

[15] Form PF is an ongoing collection, and data from later collections may be included in future analyses.

observations per fund and 44 monthly return observations per fund.

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 $30 million. The 90$^{\text{th}}$ percentile fund in TASS has just under one billion dollars in assets under management. Thus, this paper represents one of the first empirical analyses of the illiquidity characteristics of large hedge funds, which comprise the vast majority of total hedge fund assets.[16]

The bottom panel of Table 1 contains counts of funds by strategy. Our strategy categorizations are determined by the value of funds' assets reported in each of 22 pre-selected strategy categories, which fall within seven broader classifications: Credit, Equity, Event Driven, Macro, Relative Value, Managed Futures, and Other.[17] A hedge fund is categorized as pursuing a given broad strategy if 75% or more of its assets are allocated to that strategy. If there is no broad strategy category that contains 75% or more of the fund's assets, we classify the fund as multi-strategy. Table 1 shows Equity is the most common strategy, with 808 funds, while Managed Futures is the least common strategy. 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, 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 among asset illiquidity, investor share restrictions, risk-adjusted returns, and fees. In this section, we develop measures of portfolio illiquidity and investor share illiquidity based on estimated time horizons reported by funds. We then describe the properties of these illiquidity measures, and estimate the extent to which they are correlated.

[16]In concurrent work, Aragon, Ergun, Getmansky, and Girardi (2017b) and Aragon, Ergun, Getmansky, and Girardi (2017a) examine liquidity characteristics and risk management in hedge funds using Form PF data.

[17]This is question 20 on Form PF.

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."

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 measure of portfolio illiquidity, $L_{p,i,t}$, as:

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

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 horizon (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 2,550 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; 893 funds report an average value of portfolio illiquidity that is greater than 30 days, and 503 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 (excluding Other) are Event Driven and Credit, with median values of 41.3 and 26.6, 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.[18] 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 on the form 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 introduces some noise in our estimates of the most deeply illiquid assets.

The expected portfolio illiquidity reported by funds represents an estimate of an unobservable quantity, and the information content of the portfolio illiquidity measure is of first-order importance in this study. Fortunately, the data offers some reassurance that L_p does indeed capture significant economic variation. First, as discussed previously, funds with strategies one would expect *ex ante* to be associated with the most liquid assets report the shortest liquidation horizons. Second, as we show later, portfolio illiquidity and investor share illiquidity are highly correlated. This is encouraging because investor share illiquidity is directly observable from funds' offering documents. Further, the correlation between the length of funds' contractual borrowing arrangements, which is

[18] 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.

similarly observable, and portfolio illiquidity is also positive and economically large.[19] As one final piece of evidence, Form PF asks funds to report the total value of assets priced by the Generally Accepted Accounting Principles' (GAAP) *Level 1*, *Level 2*, and *Level 3* classifications.[20] *Level 1* indicates assets that have prices observable from trades, whereas *Levels 2* or *3* indicate prices that are not directly observable and instead must be approximated by a model. Presumably, infrequently traded illiquid assets are more likely have prices that fall under *Levels 2* or *3*.[21] In our data, the correlation between the within-fund average portfolio illiquidity $\overline{L}_p$, and the within-fund average of the fraction of assets that are *Level 1* is -44.69%; the correlation with the average fraction of assets that are *Level 3* is 77.63%. Together, these facts suggest that while L_p is likely measured with error, it nonetheless appears to be an economically meaningful estimate of the illiquidity of fund 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 for portfolio illiquidity. As with portfolio illiquidity, we construct a measure of investor share illiquidity, $L_{s,i,t}$, for each fund i and quarter t:

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

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, it is assumed that $\overline{W}_7 = 366$. We also compute the

[19] Results available upon request. See also Aragon, Ergun, Getmansky, and Girardi (2017b) and Aragon, Ergun, Getmansky, and Girardi (2017a).

[20] See U.S. GAAP Codification of Accounting Standards, SFAS No. 157, September 2006: Fair Value Measurements.

[21] See Getmansky, Lo, and Makarov (2004) for a discussion of mark-to-model pricing and illiquidity.

within-fund average of $L_{s,i,t}$:

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

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 10% of fund-quarter observations indicate 100% of NAV can be redeemed within seven days), 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 (excluding Other) are found in Event Driven and Credit funds, which have a median investor share illiquidity of 283.6 and 250.8 days, respectively.

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 arise 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. The coefficient value of 0.498 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 fund 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 estimate panel regressions of the relationship between $L_{p,i,t}$ and $L_{s,i,t}$ and include time and fund fixed effects separately. The results from the specification with time fixed effects show a coefficient on L_s that is highly similar to the one estimated in the cross-sectional regressions. In the specification with fund fixed effects, however, the coefficient is dramatically reduced, and is only significant at the 5% 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.[22] 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.

We reiterate that the strong, reduced-form relationship between portfolio illiquidity and investor share illiquidity is not assured. Rather, it is only one of many possible equilibrium outcomes, and the size of the correlation between L_p and L_s has implications for the allocation of illiquidity costs (and benefits) between fund managers and investors. 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 price impact from forced asset sales is passed from the manager to fund investors. Such an outcome would make sense if

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

hedge fund investors could more easily diversify their investments across various asset classes. Given the large investments of personal wealth in hedge funds by fund managers, we view this as a reasonable possibility.

We interpret the tight relationship between portfolio illiquidity and investor share illiquidity as evidence that fund managers pass at least some of the risks associated with illiquid assets to fund investors. However, Table 2 and Figure 2 show that on average portfolio assets are considerably more liquid than investor shares. Does this imply that *all* of the illiquidity risks are passed to fund investors? Not necessarily. Recall that the portfolio illiquidity measure is based on expected liquidation horizons during normal market conditions. If an increase in the average level of asset illiquidity is associated with higher price impact risk during times of market stress, or greater uncertainty in the realized liquidation period, then a proportional increase in share restrictions may not be enough to account for the entire increase in illiquidity *risk* to the manager, even if L_s increases faster L_p.

Additional evidence in support of this hypothesis is prevalent. For example, the vast majority of funds in our data ($> 85\%$) maintain the right to impose gates or suspensions on redemptions for 100% of fund NAV. These extraordinary measures are in addition to funds' normal lockups and notice periods, and if share restrictions were sufficient to eliminate the entirety of redemption risk, the option to impose such measures would not be necessary. Moreover, Table 3 shows that share illiquidity on average increases at almost double the rate of portfolio illiquidity. If the risks associated with a one-day increase in L_p could be fully covered by a one-day increase in L_s, the growth in L_s should be equal to the growth in L_p. In fact, material restrictions such as gates appear to be positively related to liquidity buffers; a regression of the fraction of net assets that can be materially restricted (gated) on the value $L_s - L_p$ produces a coefficient of 0.049 (0.049%) with a t-statistic of 4.68. That is, funds with a larger buffer between investor share illiquidity and portfolio illiquidity also maintain the ability to restrict redemptions on a larger portion of the fund's net asset value. The association between $L_s - L_p$ and the duration of funds' contractual borrowing arrangements is also positive, suggesting funds that increase investor redemption buffers also increase financing liquidity buffers. In total, these results suggest that while L_p well-represents the level of portfolio illiquidity, it does not fully capture the entirety of illiquidity risk to the hedge fund manager. In the following section, we investigate this further by examining the fraction of the illiquidity premium captured by fund managers. If share restrictions are sufficient to shield the manager from all of the illiquidity risk associated with price impact from forced asset sales, then the entire illiquidity premium should be passed to fund investors.

5 The Illiquidity Premium and Hedge Fund Performance

In this section, we estimate the size of the illiquidity premium from the illiquidity of hedge funds' portfolios, and evaluate the extent to which illiquidity exposure is an important component of hedge fund returns.

The size of the illiquidity premium is a central issue in the study of systemic risk. An argument based on revealed preference suggests that if the illiquidity premium is large, market participants must view illiquidity as sufficiently costly or risky to demand significant compensation for bearing it. Alternatively, a small illiquidity premium would suggest that market participants require relatively little reward for holding illiquid assets, and must therefore view illiquidity risks as less important.

The reliance of funds' returns on illiquidity exposure also has important systemic risk implications, as it indirectly proxies for the extent to which funds may be exposed to illiquidity risk. If funds rely heavily on illiquid assets to produce competitive returns, then they may be materially exposed to illiquidity risk. Brunnermeier and Pedersen (2009) show that the forced unwinding of illiquid positions can lead to "liquidity spirals" that propagate risks across the financial system. Alternatively, if funds' returns are not reliant on illiquidity, then they may be protected from fire sale risks and better able to weather economic downturns.

5.1 A Baseline Model without Illiquidity

Our estimation strategy to measure the size of the illiquidity premium and its importance in the cross-section of hedge fund returns follows the standard two-stage regression approach of Fama and MacBeth (1973). In the first stage, we estimate betas on factors from time-series regressions run fund-by-fund:

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

where $R_{i,t}$ is the monthly rate of return for fund i in period t, 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 economic exposures in hedge fund returns (Fung and Hsieh (2001), Fung and Hsieh (2004))[23]. The FH model includes portfolios representing stock market exposures, interest rates, credit spreads, and trend-following factors for currency, commodity, and bond markets.[24] Finally, our third model combines the seven FH factors with the three Pastor and Stambaugh (2003) liquidity measures (FH+PS). The goal of the FH+PS model is to ensure that the L_p and L_s measures of illiquidity are orthogonal to factors based on time-varying, systematic market liquidity.[25]

In much of our analysis, the (FH+PS) model will be our primary specification, but the 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. Each is based on gross-of-fee returns. The intercepts in equation (9) represent one measure of hedge fund alphas (Jensen's alpha). Overall, the data show that hedge funds earn 70 basis points per month on average, with Equity being the most lucrative strategy at 86 basis points per month, followed by Multi-strategy and Event Driven, each earning 66 and 67 basis points per month on average, respectively. Macro had the worst performance at 40 basis points per month. Funds also show substantial Jensen's alphas; the average monthly alpha almost always exceeds 20 basis points, and averages 35–51 basis points depending on the model. The largest alphas

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

[24] Specifically, the seven Fung-Hsieh style portfolios 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.

[25] Strictly speaking, only one of the three Pastor and Stambaugh (2003) measures is a traded risk factor. We include all three measures to control for residual variation in market-wide liquidity that may not be reflected in the traded factor.

come from Relative Value and Credit strategies, and the lowest come from the Macro and Equity strategies. 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.

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 the level of illiquidity in fund portfolios:

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

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

$$\widehat{\phi} = \frac{1}{T} \sum_t \widehat{\phi}_t, \quad \widehat{\alpha}_i = \frac{1}{T} \sum_t (\widehat{\phi}_{0,t} + \widehat{a}_{i,t}), \tag{12}$$

where (10) includes an intercept ($\phi_{0,t}$) 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: $\frac{1}{T}\sum_t(\widehat{\phi}_{0,t} + \widehat{a}_{i,t})$.

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 41%, with a high value of 84%, 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 hedge fund returns. 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. This suggests that during our sample period, exposure to systematic liquidity risk is actually associated with lower expected returns for hedge funds.

Table 6 shows the correlations between $\widehat{\beta}_i$ from the FH+PS first-stage regressions, and the averages of log portfolio and investor share illiquidity, $\overline{\log}(L_p)$ and $\overline{\log}(L_s)$. The largest correlations are with the bond-fund factor (β_{bd_mf}) at 23% and 31%, respectively. The largest negative correlations are with the equity market factor, at -22% and -11%, respectively. The correlations between our measures of illiquidity and the loadings on the Pastor-Stambaugh liquidity measures 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.

5.2 The Illiquidity Return Premium

Despite its importance for inferring the risks associated with illiquid assets, estimating the size of the illiquidity poses various challenges. 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)) and short-term reversals (Pastor and Stambaugh (2003)). But assets that trade enough for such measures to be accurately 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 standard liquidity measures to be well-defined.

The data described in Sections 3 and 4 offer a way forward. Instead of estimating the illiquidity of individual assets and calculating their expected returns, we can 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.[26]

We obtain estimates of the size of the illiquidity premium by re-estimating (10) with the addi-

[26]The same motivation is offered in Khandani and Lo (2011), who estimate the illiquidity premium from hedge fund portfolios based on measures of return smoothing as a proxy for asset illiquidity.

tion of hedge fund portfolio or investor share illiquidity as an explanatory fund characteristic:

$$R_{i,t} = \lambda_t^{\ell} \log(L_{i,\ell,t}) + \gamma_t' \beta_i + (a_{i,t}^{\ell}), \tag{13}$$

$$E[R_{i,t}] = \widehat{\lambda}^{\ell} E_i[\log(L_{i,\ell,t})] + \widehat{\gamma}' \beta_i + \widehat{\alpha}_i^{\ell}, \tag{14}$$

$$\widehat{\gamma} = \sum_t \widehat{\gamma}_t / T, \;\; \widehat{\alpha}_i^{\ell} = \frac{1}{T} \sum_t (\widehat{\phi}_{0,t}^{\ell} + \widehat{a}_{i,t}^{\ell}), \;\; \widehat{\lambda}^{\ell} = \sum_t \widehat{\lambda}_t^{\ell} / T, \tag{15}$$

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 $\ell = p, s$.

Estimating equation (14) delivers two coefficients of interest. The first is $\widehat{\lambda}^{\ell}$, which measures the return premium associated with portfolio or investor share illiquidity. The second is $\widehat{\alpha}_i^{\ell}$, which is average risk-adjusted return for fund i accounting for the exposure to illiquidity.

5.2.1 Portfolio Illiquidity

We first estimate the return premium implied by hedge funds' investments in illiquid assets. The parameter $\widehat{\lambda}^p$ in equation (14) gives the additional, *ceteris paribus* expected return earned from an additional log-day of portfolio illiquidity, $\log(L_p)$. Thus, $\widehat{\lambda}^p$ has a natural economic interpretation — the increase in return associated with the additional time needed to sell assets without price impact. Illiquidity premia estimated from return serial correlation or other such measures do not have this straightforward interpretation.

We obtain a coefficient estimate of $\widehat{\lambda}_{gross}^p = 0.0467$ for the illiquidity premium estimated from gross-of-fee returns. 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. $\widehat{\lambda}_{gross}^p$ implies that an additional log-day of illiquidity is associated with an additional 0.047 basis points of returns per month, or 56 basis points per year. The parameter $\widehat{\lambda}_{gross}^p$ implies that a fund in the 5^{th} percentile of illiquidity will, all else equal, earn 54 basis points per month in risk-adjusted returns on average across all strategies. By comparison, a fund in the 95^{th} percentile will earn over 82 basis points a month. This gives a gross illiquidity premium of nearly 28 basis points per month, equivalent to an annual return premium of over 3.4% per year.

5.2.2 Investor Share Illiquidity

We next examine the return premium associated with investor share illiquidity, which likely represents excess returns from both the illiquidity of the assets and from greater manager investment discretion and flexibility, such as ability to pursue risky arbitrage or investments with uncertain payoff horizons. To do so, we re-estimate equations (13) and (14) but now include log investor share illiquidity, $\log(L_s)$, as the measure of illiquidity on the right-hand side.

The premium estimated from gross returns is $\widehat{\lambda}^s_{gross} = 0.0632$, which indicates that an additional log-day of investor share restrictions increases fund performance by over six basis points per month, or an additional 76 basis points per year. The premium estimated from investor share illiquidity is over 35% higher than the premium estimated from portfolio illiquidity, consistent with tighter investor share restrictions providing return benefits beyond their association with more illiquid assets.

Of course, because investor share restrictions embed portfolio illiquidity (as demonstrated by their strong correlation), it is difficult to interpret the premium earned by investor share illiquidity. Instead, a more informative result may be the investor share illiquidity return premium earned *conditional on the illiquidity of the assets*. This provides a measure of the additional return benefits associated with more tightly restricted shares beyond those imposed to match the illiquidity of the assets. To examine this, we estimate the second-stage model including both portfolio illiquidity and investor share illiquidity on the right hand side:

$$R_{i,t} = \lambda_t^{s|p} \log(L_{i,s,t}) + \lambda_t^{p|s} \log(L_{i,p,t}) + \gamma_t' \beta_i + (a_{i,t}^{p,s}), \tag{16}$$

where the superscripts $p|s$ and $s|p$ denote portfolio illiquidity conditional on investor share illiquidity, and investor share illiquidity conditional on portfolio illiquidity, respectively. In equation (16), the coefficient $\lambda_t^{s|p}$ now has the interpretation of the additional, risk-adjusted return earned by tighter share restrictions *controlling* for the illiquidity of the assets.

We obtain parameter estimates of $\widehat{\lambda}^{s|p}_{gross} = 0.1198$ and $\widehat{\lambda}^{p|s}_{gross} = 0.0754$. The premium earned by investor share illiquidity conditional on the illiquidity of the assets is much larger than the estimate derived from the model without portfolio illiquidity (0.1198 vs. 0.0632). This may suggest substantial return benefits associated with share restrictions that result independent of asset illiquidity.

Our finding that investor share illiquidity delivers greater risk-adjusted returns than portfolio illiquidity is not particularly surprising. In addition to matching the illiquidity of assets and liabilities, previous research has found that tighter share restrictions may facilitate higher returns through greater managerial discretion (Agarwal, Daniel, and Naik (2009)), through the ability to exploit opportunistic trading opportunities during times of market stress and tight funding conditions (Shleifer and Vishny (1997), Aragon, Martin, and Shi (2019)), or through profitable trades with uncertain payoff horizons (Giannetti and Kahraman (2017)). Thus, we should expect the excess returns associated with tighter share restrictions to comprise both the illiquidity premium earned on funds' illiquid assets as well as additional returns that arise from greater trading flexibility and discretion. Additionally, our measure of portfolio illiquidity only captures the expected liquidation horizon during normal market conditions. But if higher order moments of illiquidity are priced (such as the variance of the liquidation horizon), this may be reflected both in expected returns and in endogenously tighter lockups and redemption notice periods. Finally, it may simply be that more illiquid shares result from skilled managers exercising greater bargaining power, in which case investor share illiquidity will be associated with higher risk-adjusted returns simply because it is associated with greater manager skill.

We offer a word of caution about the interpretation of these findings. Specifications with investor share illiquidity do not account for potential confounding factors, such as skilled managers exercising greater bargaining power. It may be that the benefits we ascribe to tighter share restrictions simply reflect that managers with more skill can demand greater concessions from prospective investors, in which case tighter share restrictions merely proxy for manager ability. Further, as we argued in Section 4.3, a widening gap between investor share illiquidity and portfolio illiquidity could signal a *greater* illiquidity risk to the fund manager if the illiquidity buffer doesn't increase proportionally with risk; the manager capturing some of the investor share illiquidity premium conditional on the level of portfolio illiquidity offers support for this possibility. Without a credible instrument or plausible exogeneity, we are unable to rule out these alternative hypotheses, and the results should be viewed only as instructive.

5.3 Hedge Fund Alphas and Illiquidity

Our analysis of the illiquidity premium suggests that illiquid assets earn significant compensation in the form of higher risk-adjusted returns. The magnitude of the premium suggests the costs and risks of holding illiquid assets may be substantial. In this section, we investigate whether hedge funds' risk-adjusted returns (alphas) are significantly dependent on the illiquidity premium.

The top panel of Figure 4 plots monthly returns against the average of the log of portfolio and investor share illiquidity, respectively. Figure 4 also plots two sets of estimated alphas based on the FH+PS factor model: those estimated as the intercept in (9) ($\beta_{0,i}^{FH+PS}$) and those estimated from the second-stage regressions (α_i^{FH+PS}). In each panel, funds are sorted into 20 groups based on average log portfolio and investor share illiquidity, $\overline{\log}(L_p)$ or $\overline{\log}(L_s)$, respectively. The returns plotted on the vertical axis are averages within each group.

The results in Figure 4 show that monthly gross-of-fee returns and alphas estimated from standard factor models demonstrate a clear and strong relationship with hedge fund illiquidity. The return differences implied by Figure 4 between the most and least liquid portfolios are economically large, ranging from 40 to almost 80 basis-points *per month*. For comparison, the median monthly gross return in our sample is 68 basis points. Figure 4 offers compelling evidence that portfolio and investor share illiquidity are important characteristics of cross-section of hedge fund returns.

While Figure 4 offers visual evidence that 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 $\log(L_{i,p,t})$ 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 the GRS statistic:

$$\frac{T-N-K}{N}\left(1+\bar{f}'\hat{\Omega}^{-1}\bar{f}\right)^{-1}\hat{\alpha}'\hat{\Sigma}^{-1}\hat{\alpha} \sim F_{N,T-N-K}, \tag{17}$$

where $N=5$ is the number of portfolios, $K=10$ is number of factors, $\bar{f}$ 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 15, which implies a p-value of zero to the 8^{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.

If hedge fund illiquidity is an important component of hedge fund returns, we should see that alphas estimated from equation (14), $\widehat{\alpha}_i^{\ell}$, are considerably smaller than alphas estimated from equation (10). 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 gross-of-fee 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 the level of illiquidity. We focus on gross alphas because we are interested in the effect of illiquidity on the returns to fund assets. We later explore the net-of-fee returns earned by fund investors. The results in Table 7 confirm that illiquidity is an integral component of hedge fund returns. In the FF model, portfolio illiquidity reduces average gross alpha from 0.51% per month to 0.34% per month; portfolio illiquidity can explain about 33% of average alpha estimated from the Fama-French-Carhart model. Portfolio illiquidity reduces the mean absolute pricing error by around 22%. Conversely, investor share illiquidity reduces average alpha in the FF model by 67%, and mean absolute pricing errors by 36%.

In the FH model, portfolio illiquidity reduces average gross alpha from 0.55% to 0.41% (a 25% decrease) and mean absolute pricing errors from 0.59% to 0.48% (a 19% decrease). Once again, investor share illiquidity does a better job explaining average alpha and absolute pricing errors than portfolio illiquidity, with reductions of 51% and 36%, respectively. The results from the FH+PS model are highly similar; portfolio illiquidity reduces average alphas and mean pricing errors by 27% and 19%, respectively, and investor share illiquidity reduces them by 54% and 36%, respectively. In all three models, investor share illiquidity contains much more explanatory power for hedge fund alphas and mean absolute pricing errors than portfolio illiquidity. We expound on this result below.

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 largely similar for 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^{th} through the 90^{th} percentiles. This suggests the reduction in alphas is the result of a leftward shift of the distribution, rather than a reduction in skewness or an effect on higher-order moments.

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). We use the parameter θ_0, which reflects the fraction of the true economic return comprised by the contemporaneous reported return; higher values of θ_0 are intended to proxy for more liquid portfolios. The inclusion of θ_0 appears to have no effect on the distribution of alphas or mean pricing errors. 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 may exhibit larger estimated alphas than samples with smaller funds.[27] Finally, it could be that the funds that choose not to report to public databases are those with bespoke strategies that carry the greatest threat 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.

[27] 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)).

5.4 An Illiquidity Risk Factor

The previous sections have established that illiquid funds earn considerably higher returns on average than liquid funds, and that fund illiquidity can explain between 27-55% 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 liquidity risk factor.

This distinction is important for understanding the financial stability implications of hedge fund illiquidity. If the illiquidity premium results from exposure to systematic illiquidity risk, this means funds' returns are expected to be worse during times of poor market liquidity. If fund investors are more likely to redeem shares during times of bad performance, then funds may be forced to liquidate assets exactly when markets are most illiquid. Alternatively, if the illiquidity premium results from holding assets that are difficult to sell quickly regardless of market illiquidity, then funds may less exposed to fire sale risks driven by unanticipated redemptions.

To disentangle these two possibilities, 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 liquidity 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%.

5.4.1 Factor Construction

We sort funds into five groups based on their average portfolio illiquidity, $\overline{L}_p$. Because our sample includes just under five years of returns, we sort funds only once. This is likely not too problematic because portfolio illiquidity is highly persistent. In each month and year, we construct the 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 5 shows the time series distribution of the illiquidity risk factor, normalized to have mean zero and standard deviation one. The un-normalized mean value of the factor is 0.113% with a standard deviation of 0.651%. On average, the 20% most illiquid funds outperform the 20% most liquid by roughly 1.36% per year. For comparison, the top panel also includes the (normalized) traded liquidity factor from Pastor and Stambaugh (2003), which has a correlation of 32% with the illiquidity risk factor constructed here. This suggests the illiquidity factor based on L_p does capture some variation in market-wide liquidity.

The bottom panel of Figure 5 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 77% 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 1.03% per month, compared to 0.76% for the illiquid portfolio.

5.4.2 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 coefficients on the illiquidity factor betas range from -1.74 to 1.22 over the 57 monthly cross-sectional regressions, with a mean of -.041 and a standard deviation of 0.676. This results in a t-statistic of -0.458, 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.

5.4.3 Return Smoothing

The lower volatility of average returns associated with illiquid funds highlights an important consideration for the illiquidity premium: *return smoothing*. A well-known property of hedge fund returns is that stale prices due to illiquid assets give rise to autocorrelation in 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)). Indeed, the average first-order autocorrelation coefficient for the 20% most liquid funds in our data is virtually zero (0.0003), while the average for the 20% least liquid funds is over 15.3%. Further, these averages grow monotonically across each illiquidity quintile.

However, return smoothing is unlikely to be an important explanatory factor for the empirical relationships between estimated alphas and $\log(L_p)$ and $\log(L_s)$. First, as shown in the top panel of Figure 4, there is a strong empirical relationship between average raw monthly returns and the log of average portfolio illiquidity; while smoothing would affect estimated alphas by mechanically shrinking betas, it would not affect raw returns. Second, in a cross-sectional regression of estimated alphas on the within-fund average of log portfolio illiquidity, the coefficient on log illiquidity is virtually unchanged when the first-order autocorrelation coefficient, ρ_i, is included as an additional control. Further, the coefficient on ρ_i is negative and statistically insignificant, suggesting that holding portfolio illiquidity fixed, higher levels of smoothing are not associated with larger alphas.

Finally, while the lower volatility (not smoothing) of returns of the most illiquid funds will also shrink 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.

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 Who Captures the Illiquidity Premium?

Section 5.2 established that the premium earned on illiquid assets is large and has significant explanatory power for hedge fund returns. Next, we examine who ultimately captures the illiquidity premium, fund managers or fund investors. Section 2 proposed an economic framework for interpreting the pass-through rate of the illiquidity premium. If fund managers have a large personal wealth stake in the fund, then redemption and fire sale risks are difficult to hedge. If managers do not pass the entirety of these risks to fund investors through the imposition of share restrictions, then managers may require additional compensation for bearing such risks.

The illiquidity premium pass-through rate has important implications for financial stability. Because fund managers may have more difficulty hedging illiquidity risk, if managers retain a substantial portion of the risk it could exacerbate the likelihood of a rapid unwinding of illiquid positions. However, if fund investors bear the illiquidity risk, they are likely better able to diversify it and the systemic risks associated with illiquid assets in hedge fund portfolios may be partially mitigated. Ideally, one would like to measure the fraction of illiquidity *risk* retained by fund managers or passed to fund investors, but how to empirically measure such a decomposition is unclear. Instead, we argue that in competitive markets those who bear the risks of illiquid assets must be the ones who earn the premium. Thus, the pass-through rate of the illiquidity premium from managers to investors serves as a proxy for the pass-through rate of illiquidity risk.

To date, a decomposition of the fraction of the illiquidity premium retained by managers versus passed to fund investors has not been possible. The data in this paper allow for such an analysis. Further, we are able to distinguish between the return premium earned solely from the illiquidity of the assets, and the additional premium earned by share restrictions that exceed those necessary to match the liquidity of assets and liabilities. To our knowledge, each of these findings is novel to the existing literature.

6.1 The Pass-Through Rate Implied by Returns

In Section 5.2.1 we found that the return premium on portfolio illiquidity was $\widehat{\lambda}^{p}_{gross} = 0.0467$ based on gross returns. 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. The estimated value of $\widehat{\lambda}^{p}_{net}$ is 0.0362, equivalent to an annual expected excess return of 0.44% for an additional log-day of investor share illiquidity. Investors earn an illiquidity premium of around 22 basis points per month between the 5$^{\text{th}}$ and 95$^{\text{th}}$ illiquidity-percentile of funds, or around 2.7% per year.

The t-statistic on $\widehat{\lambda}^{p}_{net}$ is 3.62, which allows us to strongly reject the null that the return premium earned by fund investors is zero. The ratio $\widehat{\lambda}^{p}_{net}/\widehat{\lambda}^{p}_{gross}$ provides an estimate of the fraction of the illiquidity premium passed through to fund investors — it is the extra return investors expect from an additional log-day of portfolio illiquidity, relative to the additional return earned by the assets. The coefficient estimates imply that 77% 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. Taken together, this indicates that a substantial portion of the illiquidity premium earned by fund assets are passed through to fund investors.

Similarly, the share illiquidity premium estimated from net returns is $\widehat{\lambda}^{s}_{net} = 0.0454$, compared to an estimate of $\widehat{\lambda}^{s}_{gross} = 0.0632$ based on gross returns. This implies a pass-through rate to investors of the share illiquidity premium of $\widehat{\lambda}^{s}_{net}/\widehat{\lambda}^{s}_{gross} = 72\%$. That is, nearly three-quarters of the additional returns earned from investor share restrictions are passed through to fund investors. This suggests that the rewards earned by investors' shares being more tightly restricted are passed through to fund investors at a rate nearly equal to the rate estimated from portfolio illiquidity.

Once again, we attempt to disentangle (with the appropriate caveats discussed above) the return premium earned on tighter share restrictions beyond that needed to match the illiquidity of the assets. We estimate $\widehat{\lambda}^{s|p}_{net} = 0.0929$ and $\widehat{\lambda}^{p|s}_{net} = 0.0526$. The ratio $\widehat{\lambda}^{s|p}_{net}/\widehat{\lambda}^{s|p}_{gross} \approx 78\%$ suggests that nearly 80% of the return benefits achieved by further restricting investor shares, conditional on the illiquidity of the assets, are passed to fund investors. While the entirety of the *costs* associated with increased share restrictions are borne by fund investors, and may therefore suggest a competitive pass-through rate of 100%, we note that we are unlikely to statistically reject the null of a 100% rate. Further, if share restrictions help to coordinate investor incentives by limiting the scope of first-mover advantages in redemptions (as in standard bank-run models), investors may accept lower total returns in exchange for such coordination.

6.2 The Pass-Through Rate Implied by Fees

Section 4.3 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, the previous section showed that the illiquidity premium estimated from fund assets is 22-28% higher than the premium earned by fund investors. This suggests that fund managers are capturing some of the illiquidity premium as compensation for price impact risk that arises either from potential redemptions that result from not sufficiently restricting investors' shares, or from the risk of margin calls on leveraged positions that cannot be reduced through share restrictions.

However, 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 fund managers is not very informative. 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 informative, it must be that the fees funds charge are a function of illiquidity, as 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.

The testable hypothesis is that the rates of manager 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 portfolio illiquidity or manager skill.[28] We emphasize that this analysis examines *realized* fees, which aggregate 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, and therefore also depend on funds' realized returns.

[28] 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)"

Figure 6 plots average realized fees against $\alpha_i^{p,g}$ — the alpha estimated from the FH+PS model using gross-of-fee returns and portfolio illiquidity — and against average log portfolio illiquidity, $\overline{\log}(L_{i,p})$, after controlling for $\alpha_i^{p,g}$. We add the superscript g to the alpha variable to emphasize that these are alphas estimated from gross-of-fee returns. 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* $\alpha_i^{p,g}$. 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 6, we first examine whether illiquidity mechanically generates higher realized 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 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.[29] 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 9 reports results from regressions of realized fees on manager skill, portfolio illiquidity, and other potential explanatory factors. Column (1) shows that, after controlling for fund size (net assets) and strategy, higher-alpha funds continue to earn higher fees. Column (2) shows that this relationship remains after controlling for average portfolio illiquidity; further, portfolio illiquidity predicts higher average fees even after controlling for estimated alpha. Column (3) includes the

[29]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. 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.

proportion of months in which the fund has a positive (gross) return as a proxy 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 38%, which remains highly statistically significant. Further, the coefficient on portfolio illiquidity is reduced from 0.183 to 0.126, but also remains highly statistically significant.

In column (4), we offer additional evidence that the higher fees charged by funds with more illiquid assets reflect explicit compensation for undiversifiable illiquidity risk. Regardless of the severity of investor share restrictions, one source of price impact risk managers will be unable to diversify is that arising from financing liquidity risks — the risk of forced asset sales due to margin calls or inability to rollover short-term debt. In this case, we should see more leveraged funds, who borrow more and therefore face a greater degree of financing liquidity risk, charge higher fees. Column (4) finds evidence of such an effect. Controlling for estimated skill, the fraction of positive-return months, portfolio illiquidity, and other controls, the coefficient on average fund leverage is statistically significant and economically large. A fund that is leveraged two-to-one will earn a 0.07% 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 (4) offer further evidence that the higher 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 assessed as a constant proportion of assets under management, and an incentive fee assessed as a proportion of fund returns. Unfortunately, we do not observe these separate fee components directly, and therefore must estimate them. 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 18.71%

with a mean of 16.78%, 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. 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 the Form PF data is 19.29% with a median of 19.84% and an interquartile range of 17.32-22.07%. 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 a mean of 1.47% and a median of 1.57%, and an interquartile range of 0.96-1.69%. For those that charge a 2% management fee, we estimate a mean of 1.74% and a median of 2.06%, with an interquartile range of 1.21-2.06%. In total, we view these results as strongly supportive of the credibility of our fee estimates.

To establish facts, in column (5) of Table 9 we regress the estimated incentive fee on the estimated management fee. The coefficient is positive but statistically insignificant, which suggests that the incentive fee and management fee are unrelated in our data. In column (6) of Table 9, we regress the estimated incentive fee on estimated gross-of-fee alpha, average portfolio illiquidity, average leverage, and other controls. The estimated incentive fee is economically related to both estimated alpha and portfolio illiquidity. The coefficients on each are also statistically significant, with p-values of 0.014 and 0.016, respectively. A ten basis point higher monthly alpha is associated with a 0.179% higher incentive fee. An extra log-day of portfolio illiquidity is associated with an additional 0.470% higher incentive fee. Average leverage is also economically large and statistically significant, with a fund leveraged two-to-one on average charging roughly 32 basis points higher incentive fees than an unleveraged fund. Alternatively, column (7) shows that while higher alpha is also associated with a higher management fee, portfolio illiquidity appears uncorrelated with the management fee. Leverage is weakly related to the management fee, but is marginally statistically significant and economically small.

The association between illiquidity and realized fees, and the incentive fee in particular, is consistent with models of intermediary-based asset pricing (He and Krishnamurthy (2012)) and our hypothesis that managers retain some of the illiquidity risk in fund portfolios. 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 once again that compensation for financing liquidity risk is also relevant.

In columns (8) and (9) of Table 9, we offer two more explicit tests of our interpretation that managers are compensated for illiquidity risk through higher fees. Our most compelling finding is in column (8). First, using data from Form ADV, we construct a sample of funds that are 0% manager owned. In such funds managers have no personal wealth stake in the fund. If the empirical relationship between illiquidity and fees is driven by managers having undiversifiable exposure to illiquidity, then fees and illiquidity should be uncorrelated in 0% manager-owned funds.

Indeed, that is exactly what we find. Column (8) shows that for the sample of funds with no management ownership, illiquidity is unrelated to fees. The coefficient is both economically small and statistically insignificant. This result is unlikely to be driven by lower-powered tests. Both the estimated coefficients on alpha and on leverage remain highly statistically significant, and have coefficients that are roughly equal to those estimated on the full sample (column (4)).

In column (9), we build on this intuition by examining multi-strategy funds. If managers require additional compensation due to large personal stakes in an undiversified portfolio, then more diverse strategies should reduce the potential risks of that personal commitment. That is, managers of multi-strategy funds should require less compensation for undiversifiable illiquidity risk than managers of more narrow or specific strategies. In column (9), we include an interaction between portfolio illiquidity and an indicator variable for whether the fund is multi-strategy (the coefficient on the multi-strategy dummy is unreported, but is included in the strategy fixed effects controls). In this specification, the coefficient on illiquidity is 0.160, but is highly attenuated for multi-strategy funds; the coefficient on the multi-strategy illiquidity interaction term is -.108, and is significant at the 10% level (p-value of .093). This implies that the relationship between illiquidity and fees shrinks by over two-thirds for multi-strategy funds, consistent with managers of more diversified strategies requiring less compensation for illiquidity risk.

Lastly, we note that in column (8) the coefficient on leverage remains statistically significant and similar in magnitude to the estimate in the full sample, which suggests that the association between leverage and fees is likely not due to compensation for financing liquidity risk. We view this as an interesting possible avenue for future research.

This strong attenuation of the relationship between illiquidity and fees for 0% manager owned and multi-strategy funds offers compelling evidence that managers require additional compensation for undiversifiable illiquidity risk. These findings also speak to the potential liquidity mismatch in hedge fund portfolios. Section 4 showed that in aggregate, investor shares are always more illiquid than the underlying assets. One interpretation of this finding is that there is no liquidity mismatch between hedge fund assets and liabilities. The results in Table 9, however, suggest this may not be the case. If managers have no fear of forced asset sales due to evaporating financing, then managers would require no additional compensation for holding illiquid assets. Said differently, managers will *only* charge higher fees for managing illiquid assets if there is at least one state of the world in which assets would need to be sold too quickly and managers bear the costs of selling the illiquid assets at a depressed prices.

The analysis in this section implicitly assumes that as portfolio illiquidity increases, the manager's risk also increases. Yet this may not be the case if the manager also sufficiently increases share restrictions (see the discussion at the end of Section 4.3). The ideal experiment would be to increase portfolio illiquidity, *holding investor share restrictions constant*. Unfortunately, such an analysis is not possible here because both portfolio illiquidity and share restrictions are endogenously set by managers. Further, higher fees and tighter share restrictions may be positively related because of unobserved manager skill and bargaining power, making coefficients difficult to interpret. The association between portfolio illiquidity and fees is less likely to be contaminated by manager ability, because it is less clear that managers of illiquid assets are necessarily more skilled. Finally, the results in Section 5.2 suggest that not all of the illiquidity risk is passed to fund investors; if so, increased portfolio illiquidity would imply a greater undiversifiable risk for the fund manager. Regardless of these shortcomings, when we include investor share illiquidity in our main specifications along with portfolio illiquidity, we continue to find a positive and statistically significant relationship between portfolio illiquidity and fees.[30]

[30] We also find a positive relationship between investor share illiquidity and fees, even controlling for manager skill, portfolio illiquidity, and the other controls in columns (1)–(5). This may suggest that unobserved manager skill and bargaining power are in part responsible for the magnitude of investor share illiquidity.

7 Conclusion

Illiquid assets become difficult to sell during periods of market stress, leading to the possibility of fire sales, liquidity spirals, and contagion to counterparties and unrelated markets. Because the owners of illiquid assets may bear substantial costs when they unwind their positions, they will require higher returns as compensation. Yet, despite widespread acceptance that illiquidity commands a premium, no estimates of the size of the premium exist that span the full spectrum of illiquidity. In the context of asset management, even less is known about who ultimately bears the risks associated with illiquid assets, the under-diversified intermediary or the better diversified fund investors.

We use information from the first systematic, regulatory data collection on large hedge funds to investigate these issues. In particular, we produce the first direct estimate the size of the illiquidity premium — which serves as a proxy for the potential size of illiquidity risk — based on hedge funds' self-reported portfolio illiquidity. We also estimate the illiquidity premium's importance for explaining hedge funds' risk-adjusted returns. If the illiquidity premium is a major source of hedge fund returns, this may suggest funds are significantly exposed to illiquidity risk. Finally, we estimate how much of the illiquidity premium, and therefore how much of the illiquidity risk, is ultimately captured by fund investors versus retained by fund managers.

The economic framework that gives coherence to our empirical results is 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, charge higher fees, or both. Consistent with this interpretation, we first document that investor share illiquidity and portfolio illiquidity are strongly positively correlated. This suggests that at least some of the costs of illiquid assets are passed to fund investors.

We estimate a sizeable illiquidity premium; the *ceteris paribus* return differential between funds in the 5^{th} and 95^{th} percentiles of illiquidity is nearly 3.4% per year. The large magnitude of the premium indicates that illiquidity poses a significant risk to financial market participants that in turn requires substantial compensation. Premiums estimated from alternative measures of illiquidity, such as return smoothing, fail to produce statistically significant results.

Next, we find that between 27-55% of hedge fund alpha can be explained by portfolio illiquidity, suggesting that funds rely on illiquidity as a important component of returns, and which may also indicate high exposures to illiquidity risk. We also find, consistent with greater manager discretion allowing for more profitable trading strategies, that investor share illiquidity has much greater explanatory power for hedge fund alphas than portfolio illiquidity alone.

While we find that the illiquidity premium is large and an important source of hedge fund alpha, we also find that roughly 77% of the portfolio illiquidity premium is passed to fund investors. We find a similar proportion of the premium associated with investor share illiquidity is passed to investors after controlling for the illiquidity of the assets. This suggests that much, but not all of the costs of illiquid assets are passed to fund investors. Consistent with managers retaining some of the undiversifiable illiquidity risk, 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. As predicted by theories of intermediary-based asset pricing, we also show that all of the association between illiquidity and fees arises through the performance incentive fee.

Lastly, we show that the relationship between illiquidity and fees completely disappears for funds that are 0% manager owned. That is, funds in which managers have no personal stake, and therefore have no undiversifiable exposure to illiquidity, exhibit no association between illiquidity and fees.

Our results suggest that the costs and risks associated with illiquid assets are substantial, and that this risk is largely passed to fund investors. Because fund investors are likely better able than fund managers to diversify their investments across many asset classes and strategies, the transmission of illiquidity risk from fund managers to investors likely reflects a more efficient allocation of illiquidity risk to those for whom the risks are least costly. Relative to the managers retaining all of the illiquidity risk, this allocation may also promote financial stability; if illiquidity risks are better hedged when held by investors, the impetus to extract value that results from the implicit first-mover advantage in managed funds — and which may serve as a potential catalyst for fire sales — may be mitigated in the event of deteriorating liquidity.

References

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

AGARWAL, V., N. D. DANIEL, AND N. Y. NAIK (2009): "Role of Managerial Incentives and Discretion in Hedge Fund Performance," *The Journal of Finance*, 64(5), 2221–2256.

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.

ALLEN, F., AND D. GALE (2000): "Financial Contagion," *Journal of Political Economy*, 108(1), 1–33.

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., J. S. MARTIN, AND Z. SHI (2019): "Who Benefits in a Crisis? Evidence from Hedge Fund Stock and Option Holdings," *Journal of Financial Economics*, 131(2), 345–361.

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.

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

BERNARDO, A. E., AND I. WELCH (2004): "Liquidity and Financial Market Runs," *The Quarterly Journal of Economics*, 119(1), 135–158.

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.

BRUNNERMEIER, M. K., AND L. PEDERSEN (2009): "Market Liquidity and Funding Liquidity," *Review of Financial Studies*, 22(6), 2201–2238.

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,” *The Review of Financial Studies*.

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.

He, Z., and A. Krishnamurthy (2013): “Intermediary Asset Pricing,” *American Economic Review*, 103(2), 732–70.

Jotikasthira, C., C. Lundblad, 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. (2019): "Limited Investment Capital and Credit Spreads," *Journal of Finance*, 74(5), 2303–2347.

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.**	25^{th} **Percentile**	50^{th} **Percentile**	75^{th} **Percentile**	**Obs.**
NAV ($Billions)	1.14	2.21	0.17	0.51	1.16	2,550
GAV ($Billions)	2.06	7.13	0.25	0.69	1.74	2,550
GAV /NAV	3.46	38.78	1.01	1.17	1.68	2,550
Monthly Obs./Fund	35.76	20.99	16.00	44.00	57.00	2,550
Quarterly Obs./Fund	10.64	6.39	5.00	10.00	17.00	2,550
Strategy						
Equity						808
Multi-strategy						328
Credit						227
Relative Value						217
Event Driven						191
Macro						161
Managed Futures						58
Other						560

Table 1 reports summary statistics for our main sample of 2,550 funds. Each statistic reported is the average value for the fund over the full set of observations for which we observe the fund.

Table 2: Portfolio and Investor Share illiquidity

		Portfolio Liquidity ($\overline{L}_p$)			Investor Liquidity ($\overline{L}_s$)		
	Pct. Obs.	**Mean**	**Median**	**SD.**	**Mean**	**Median**	**SD.**
All	100.0	64.5	12.1	107.0	165.1	135.5	137.6
Equity	31.7	23.9	6.3	52.8	123.0	76.2	106.7
Multi-strategy	12.9	59.2	23.0	88.9	173.2	171.9	129.4
Event Driven	7.5	91.2	41.3	106.9	254.5	283.6	115.3
Relative Value	8.5	44.8	12.4	74.8	155.0	135.5	130.5
Credit	8.9	105.4	26.6	127.9	220.6	250.8	139.6
Macro	6.3	7.0	2.9	12.6	75.2	60.5	80.3
Managed Futures	2.3	4.3	1.3	16.3	26.3	4.5	54.9
Other	21.1	135.6	55.7	147.3	218.5	315.3	159.6

Table 2 reports summary statistics for the averages of portfolio illiquidity ($\overline{L}_p$), and investor share illiquidity ($\overline{L}_s$). This table includes only the 2,550 funds in our main sample. Data reported are in days. Stars are used to replace data that fail to meet the appropriate data aggregation criteria.

Table 3: Measures of Liquidity

	(1) $\overline{L}_p$	(2) $\overline{L}_p$	(3) $\log \overline{L}_p$	(4) L_p	(5) L_p
$\overline{L}_s$	0.498***	0.452***	0.009***		
	(0.012)	(0.012)	(0.000)		
L_s				0.438***	0.045**
				(0.016)	(0.019)
Constant	-17.812***	6.570	1.397***	2.620	50.373***
	(2.541)	(5.874)	(0.103)	(7.683)	(4.978)
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	0.468	0.527	0.446	0.007
Observations	2550	2550	2550	26972	26972

Table 3 reports regression results of portfolio illiquidity L_p on investor share illiquidity 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. 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. $R_{i,t}$	0.70	0.11	0.40	0.68	0.98	1.32	0.50
	$\widehat{\beta}_{0,i}^{FF}$	0.35	-0.22	0.05	0.33	0.64	0.96	0.48
	$\widehat{\beta}_{0,i}^{FH}$	0.41	-0.17	0.07	0.41	0.72	1.04	0.49
	$\widehat{\beta}_{0,i}^{FH+PS}$	0.51	-0.13	0.17	0.51	0.83	1.19	0.54
Equity	Avg. $R_{i,t}$	0.86	0.27	0.56	0.83	1.14	1.47	0.53
	$\widehat{\beta}_{0,i}^{FF}$	0.24	-0.33	-0.07	0.21	0.52	0.86	0.50
	$\widehat{\beta}_{0,i}^{FH}$	0.26	-0.32	-0.08	0.24	0.59	0.89	0.49
	$\widehat{\beta}_{0,i}^{FH+PS}$	0.42	-0.24	0.07	0.39	0.76	1.05	0.54
Multi-strategy	Avg. $R_{i,t}$	0.66	0.22	0.42	0.59	0.89	1.20	0.43
	$\widehat{\beta}_{0,i}^{FF}$	0.42	-0.08	0.16	0.40	0.63	0.96	0.42
	$\widehat{\beta}_{0,i}^{FH}$	0.54	-0.03	0.31	0.54	0.73	1.06	0.42
	$\widehat{\beta}_{0,i}^{FH+PS}$	0.59	0.03	0.31	0.57	0.82	1.21	0.46
Event Driven	Avg. $R_{i,t}$	0.67	0.00	0.47	0.67	0.88	1.18	0.45
	$\widehat{\beta}_{0,i}^{FF}$	0.36	-0.18	0.05	0.25	0.65	0.86	0.49
	$\widehat{\beta}_{0,i}^{FH}$	0.52	-0.03	0.20	0.52	0.77	1.06	0.46
	$\widehat{\beta}_{0,i}^{FH+PS}$	0.66	0.00	0.35	0.68	0.94	1.13	0.51
Relative Value	Avg. $R_{i,t}$	0.60	0.07	0.31	0.62	0.85	1.03	0.41
	$\widehat{\beta}_{0,i}^{FF}$	0.56	0.09	0.27	0.54	0.85	1.01	0.39
	$\widehat{\beta}_{0,i}^{FH}$	0.63	0.21	0.35	0.59	0.93	1.11	0.41
	$\widehat{\beta}_{0,i}^{FH+PS}$	0.69	0.08	0.36	0.72	0.96	1.27	0.47
Credit	Avg. $R_{i,t}$	0.63	0.12	0.43	0.57	0.87	1.07	0.48
	$\widehat{\beta}_{0,i}^{FF}$	0.57	0.11	0.29	0.58	0.84	1.14	0.44
	$\widehat{\beta}_{0,i}^{FH}$	0.67	0.12	0.40	0.67	0.99	1.19	0.41
	$\widehat{\beta}_{0,i}^{FH+PS}$	0.75	0.28	0.52	0.70	0.98	1.27	0.41
Macro	Avg. $R_{i,t}$	0.40	-0.06	0.20	0.39	0.62	0.77	0.40
	$\widehat{\beta}_{0,i}^{FF}$	0.18	-0.24	-0.06	0.17	0.45	0.66	0.38
	$\widehat{\beta}_{0,i}^{FH}$	0.27	-0.18	0.00	0.23	0.59	0.72	0.41
	$\widehat{\beta}_{0,i}^{FH+PS}$	0.33	*	0.04	0.33	0.64	0.88	0.48
Other	Avg. $R_{i,t}$	0.61	-0.02	0.25	0.54	1.01	1.29	0.51
	$\widehat{\beta}_{0,i}^{FF}$	0.34	-0.18	0.06	0.30	0.68	1.02	0.49
	$\widehat{\beta}_{0,i}^{FH}$	0.39	-0.14	0.02	0.33	0.70	1.18	0.51
	$\widehat{\beta}_{0,i}^{FH+PS}$	0.46	-0.13	0.09	0.36	0.77	1.31	0.61

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 is excluded due to insufficient observations.

Table 5: Parameter Values from Second Stage Model without Illiquidity

		Percentiles						
FM Parameter Estimates	**Mean**	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}	**St. Dev**	**t-stat**
$\widehat{\phi}_{eq-mf}$	0.57	-2.39	-1.30	0.89	2.42	3.93	2.88	1.49
$\widehat{\phi}_{szspr-f}$	0.45	-3.22	-1.86	0.99	2.19	3.40	2.75	1.24
$\widehat{\phi}_{bd-mf}$	0.73	-20.81	-9.08	-1.23	12.73	25.01	21.00	0.26
$\widehat{\phi}_{crsprd-f}$	0.67	-17.30	-9.18	-1.49	12.71	26.24	17.49	0.29
$\widehat{\phi}_{agg-liq}$	-0.01	-0.05	-0.03	0.00	0.02	0.03	0.04	-1.89
$\widehat{\phi}_{innov-liq}$	-0.01	-0.06	-0.04	-0.01	0.04	0.05	0.05	-1.51
$\widehat{\phi}_{liq-v}$	0.00	-0.03	-0.01	0.00	0.02	0.03	0.02	0.00
$\widehat{\phi}_{bd-tff}$	-1.78	-19.11	-13.85	-1.26	6.42	16.90	14.87	-0.90
$\widehat{\phi}_{crrcy-tff}$	-0.57	-23.85	-14.71	-4.14	8.33	25.22	20.61	-0.21
$\widehat{\phi}_{cmdty-tff}$	-1.91	-22.12	-13.67	-4.40	8.26	17.31	17.51	-0.82

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											
$\overline{\log}(L_s)$	0.52	1.00										
$\widehat{\beta}_{bd\_tff}$	-0.08	-0.06	1.00									
$\widehat{\beta}_{crrcy\_tff}$	-0.06	-0.02	-0.49	1.00								
$\widehat{\beta}_{cmdty\_tff}$	-0.04	-0.07	0.07	-0.27	1.00							
$\widehat{\beta}_{eq\_mf}$	-0.22	-0.11	-0.06	0.06	0.00	1.00						
$\widehat{\beta}_{szspr\_f}$	0.14	0.18	-0.11	-0.14	0.03	0.24	1.00					
$\widehat{\beta}_{bd\_mf}$	0.23	0.31	-0.11	0.07	-0.19	0.02	0.08	1.00				
$\widehat{\beta}_{crsprd\_f}$	-0.01	0.13	-0.37	0.20	-0.06	0.16	0.09	0.51	1.00			
$\widehat{\beta}_{agg\_liq}$	0.12	0.14	-0.13	0.13	-0.28	0.24	0.33	0.16	0.03	1.00		
$\widehat{\beta}_{innov\_liq}$	-0.11	-0.11	0.05	-0.07	0.10	-0.22	-0.22	-0.16	-0.04	-0.80	1.00	
$\widehat{\beta}_{liq\_v}$	0.04	-0.01	0.27	-0.13	-0.16	0.14	-0.23	0.06	0.02	-0.01	-0.03	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 averages of log portfolio and investor share illiquidity, $\overline{\log}(L_p)$ and $\overline{\log}(L_s)$.

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}	0.51	0.00	0.24	0.49	0.76	1.00	0.43
α_i^{s}	0.17	-0.37	-0.09	0.16	0.42	0.68	0.43
α_i^{p}	0.34	-0.18	0.09	0.33	0.59	0.84	0.42
$\lvert\alpha_i^{FF}\rvert$	0.55	0.13	0.28	0.50	0.77	1.01	0.36
$\lvert\alpha_i^{s}\rvert$	0.35	0.06	0.14	0.28	0.49	0.73	0.30
$\lvert\alpha_i^{p}\rvert$	0.43	0.08	0.19	0.37	0.61	0.85	0.33
Fung-Hsieh							
α_i^{FH}	0.55	0.06	0.30	0.54	0.80	1.06	0.42
α_i^{s}	0.27	-0.22	0.03	0.26	0.51	0.76	0.42
α_i^{p}	0.41	-0.09	0.18	0.40	0.63	0.91	0.42
$\lvert\alpha_i^{FH}\rvert$	0.59	0.15	0.32	0.55	0.80	1.06	0.37
$\lvert\alpha_i^{s}\rvert$	0.38	0.06	0.16	0.32	0.54	0.77	0.31
$\lvert\alpha_i^{p}\rvert$	0.48	0.11	0.23	0.42	0.64	0.91	0.34
Fung-Hsieh + Pastor-Stambaugh							
α_i^{FH+PS}	0.56	0.06	0.31	0.55	0.80	1.07	0.42
α_i^{s}	0.26	-0.25	0.02	0.25	0.51	0.75	0.42
α_i^{p}	0.41	-0.10	0.18	0.40	0.65	0.91	0.42
$\lvert\alpha_i^{FH+PS}\rvert$	0.59	0.15	0.33	0.56	0.80	1.07	0.37
$\lvert\alpha_i^{s}\rvert$	0.38	0.06	0.16	0.32	0.54	0.77	0.31
$\lvert\alpha_i^{p}\rvert$	0.48	0.11	0.24	0.43	0.65	0.92	0.34
Fung-Hsieh + Pastor-Stambaugh GLM Liquidity (θ_0)							
$\alpha_i^{\theta_0}$	0.57	0.11	0.35	0.56	0.81	1.06	0.41
$\lvert\alpha_i^{\theta_0}\rvert$	0.61	0.19	0.37	0.56	0.81	1.07	0.35

Table 7 reports the distribution of estimated alphas and mean absolute pricing errors from equation (14).

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	-0.04	0.19	0.39	0.61	0.82	0.39
α_i^{FH+PS} (Gross)	0.56	0.06	0.31	0.55	0.80	1.07	0.42
α_i^{s} (Net)	0.17	-0.28	-0.03	0.18	0.38	0.59	0.39
α_i^{s} (Gross)	0.26	-0.25	0.02	0.25	0.51	0.75	0.42
α_i^{p} (Net)	0.28	-0.17	0.07	0.28	0.49	0.69	0.39
α_i^{p} (Gross)	0.41	-0.10	0.18	0.40	0.65	0.91	0.42

Table 8 reports the distribution of estimated alphas and mean absolute pricing errors from equation (14), based on the FH+PS model and separately for gross and net returns.

Table 9: Alphas, Liquidity, and Fees

	Avg. Fee	Avg. Fee	Avg. Fee	Avg. Fee	Incentive Fee	Incentive Fee	Management Fee	Avg. Fee	Avg. Fee
$a_i^{p,g}$	1.315***	1.312***	0.818***	0.802***		1.793**	0.345***	0.774***	0.809***
	(0.112)	(0.110)	(0.122)	(0.121)		(0.726)	(0.100)	(0.276)	(0.121)
$\overline{\log}(L_p)$		0.183***	0.126***	0.138***		0.470**	-0.003	-0.005	0.160***
		(0.027)	(0.027)	(0.027)		(0.196)	(0.026)	(0.044)	(0.030)
$\%\ r_{i,t} > 0$			2.880***	2.806***				4.458***	2.830***
			(0.354)	(0.351)				(0.768)	(0.351)
Avg. Leverage				0.070***		0.318***	0.031*	0.079***	0.068***
				(0.017)		(0.102)	(0.016)	(0.027)	(0.017)
Man. Fee					0.091				
					(0.299)				
Multi-strat $\times\ \overline{\log}(L_p)$									-0.108*
									(0.064)
Constant	2.005***	1.402***	-0.406	-0.492*	16.668***	12.896***	1.373***	-1.225**	-0.585**
	(0.170)	(0.189)	(0.288)	(0.286)	(0.467)	(1.316)	(0.184)	(0.590)	(0.291)
Size Controls	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Strategy Controls	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Adjusted R^2	0.166	0.201	0.250	0.262	-0.001	0.040	0.017	0.395	0.264
Observations	1022	1022	1022	1022	638	638	1038	232	1022

Table 9 shows regressions of the average realized fee, estimated incentive fee, and estimated management fee on estimated alpha, portfolio illiquidity, and other controls. $\#r_{i,t} > 0$ is the proportion 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. * denotes significance at the 10% level, ** at the 5% level, and *** at the 1% level.

8.2 Figures

Figure 1: Distribution of L_p

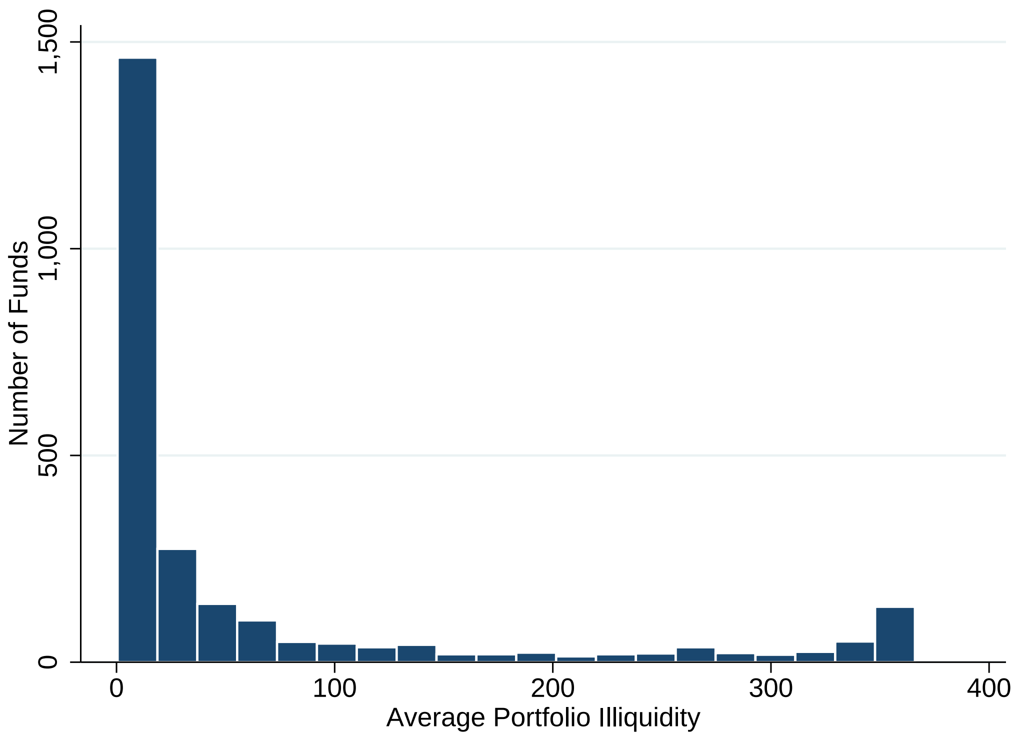

Figure 1 shows a histogram plot of the distribution of average portfolio illiquidity, $\overline{L}_p$, for our main sample of 2,550 funds.

Figure 2: Distribution of L_s

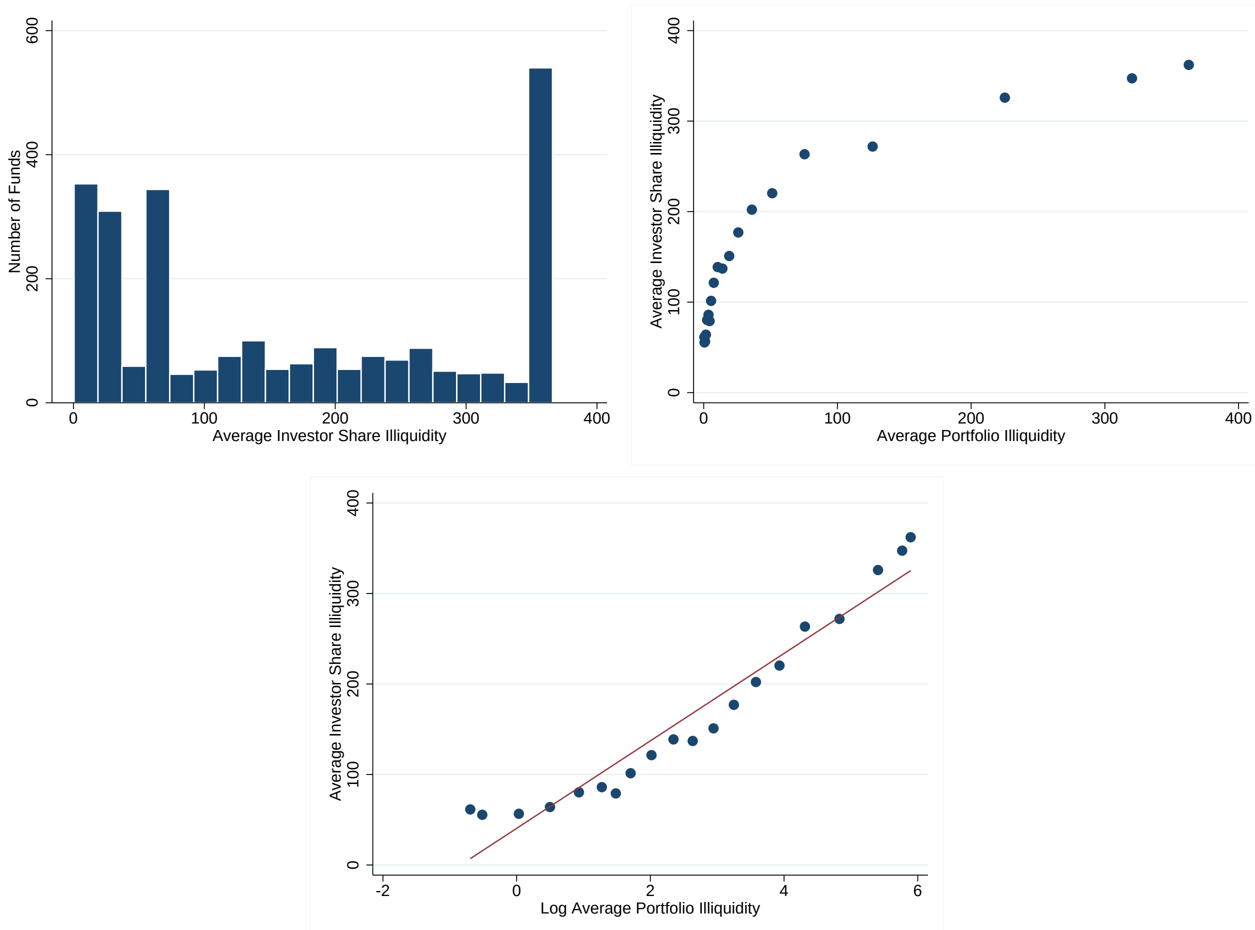

Figure 2 shows average investor share illiquidity, $\overline{L}_s$, along with average portfolio illiquidity, $\overline{L}_p$, for our main sample of 2,550 funds. The upper left panel plots the histogram of $\overline{L}_s$. The upper right panel shows a bin-scatter plot of $\overline{L}_s$ versus $\overline{L}_p$, 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)

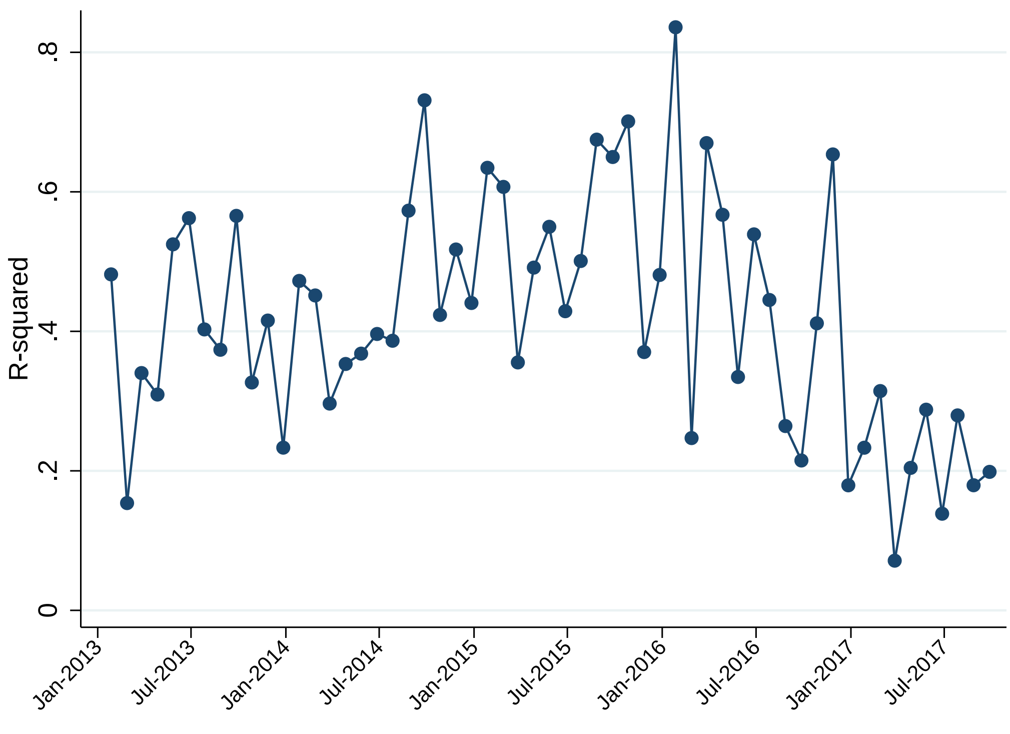

Figure 3 plots the R-squared from each of the T second-stage regressions estimated in equation (10).

Figure 4: Returns and Alphas vs. Portfolio Illiquidity

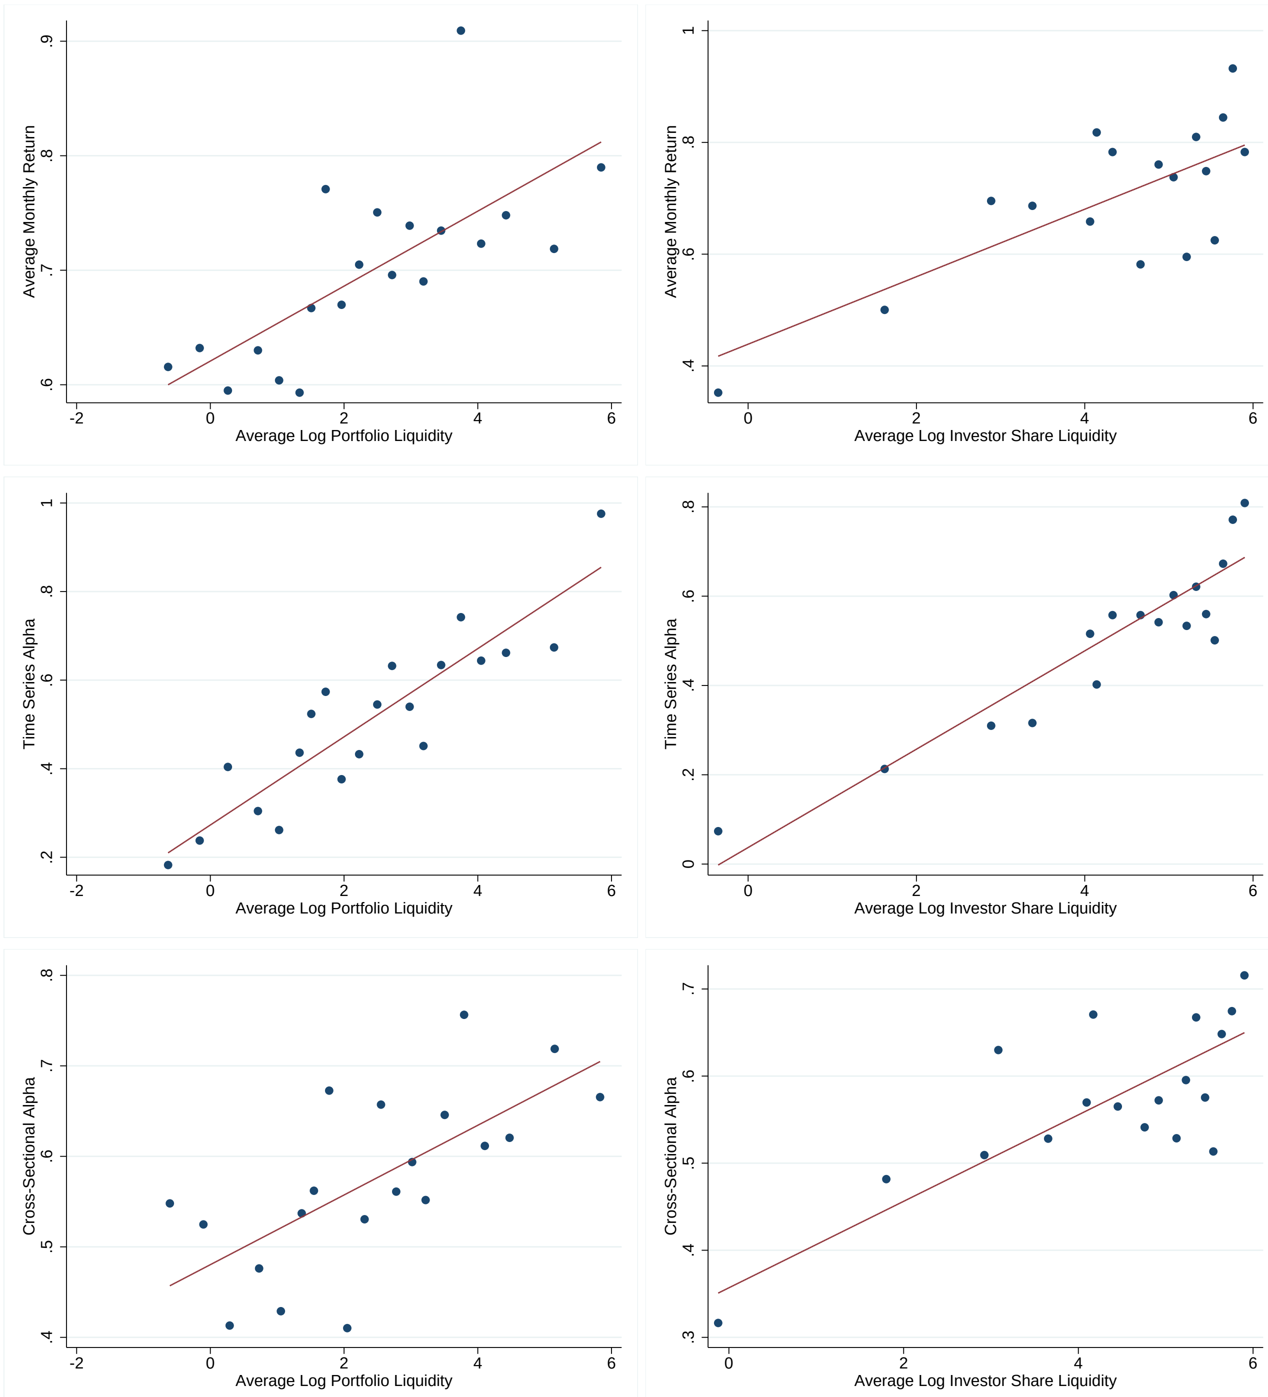

Figure 4 shows bin-scatter plots of monthly returns and estimated alphas versus average log portfolio and investor share illiquidity, $\overline{\log}(L_p)$ and $\overline{\log}(L_s)$. Time series alpha corresponds to β_0^{FH+PS} estimated in equation (9), and cross-sectional alpha refers to α_i^{FH+PS} estimated in equation (11). Funds are separated into 20 groups based on either $\overline{\log}(L_p)$ or $\overline{\log}(L_s)$, and means within each group are plotted.

Figure 5: Illiquidity Risk Factor

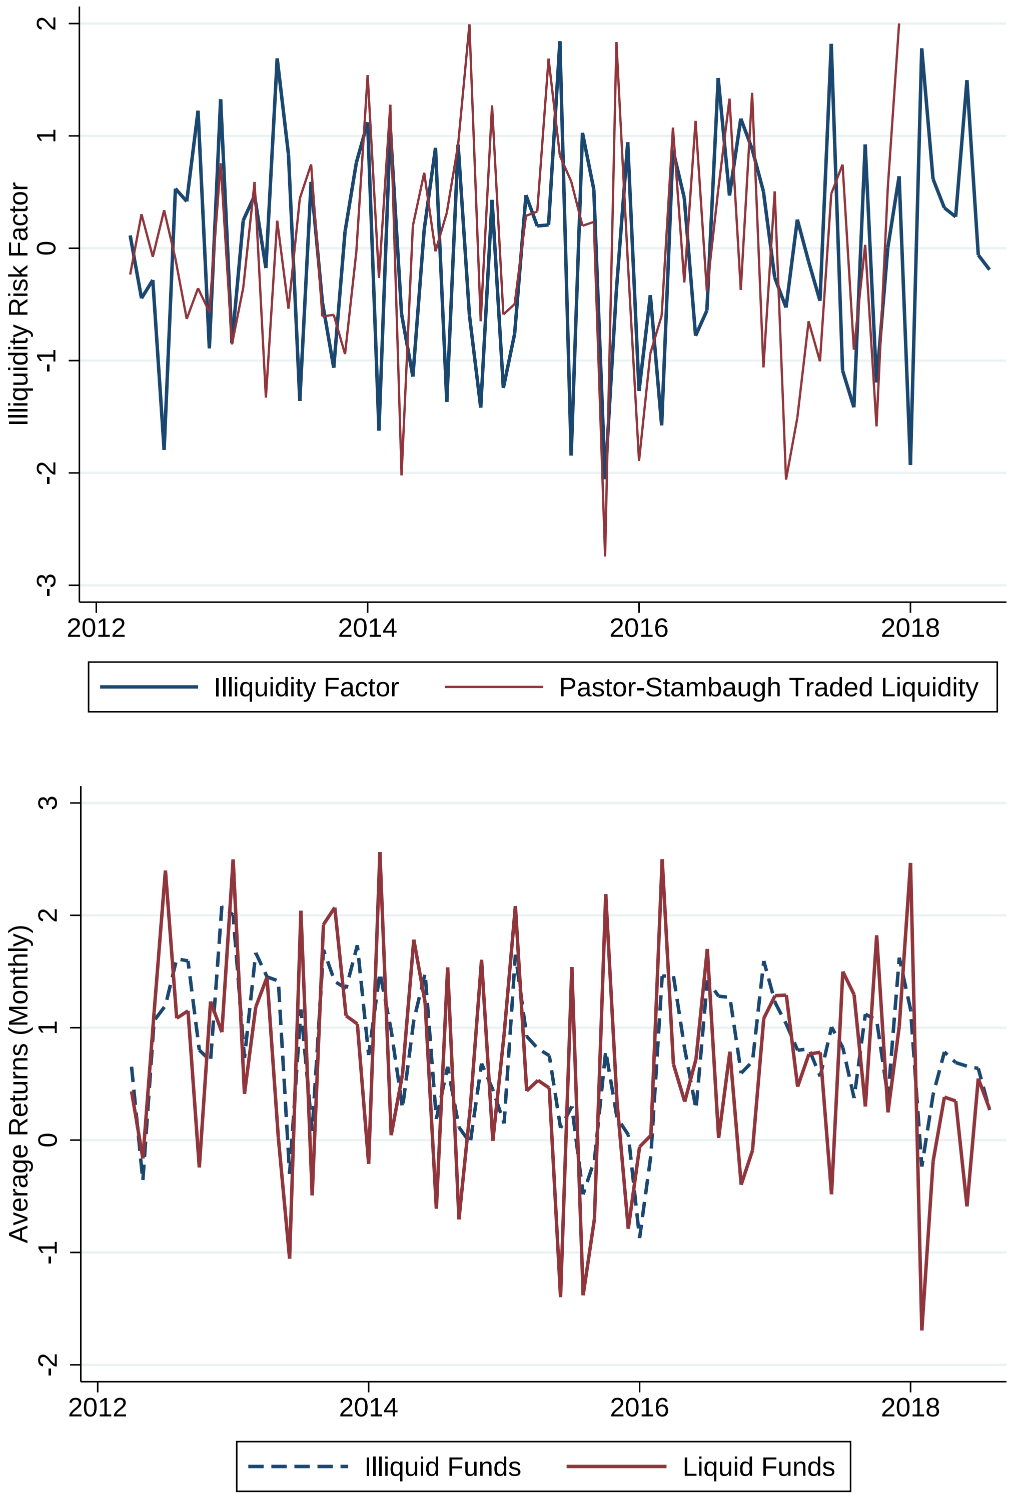

The top panel of Figure 5 shows the time-series of the illiquidity risk factor versus the traded liquidity factor from Pastor and Stambaugh (2003), each normalized to have mean zero and standard deviation one. The bottom panel shows the time series of average monthly returns for the 20% most and least liquid funds.

Figure 6: Fees vs. Alpha and Portfolio Illiquidity

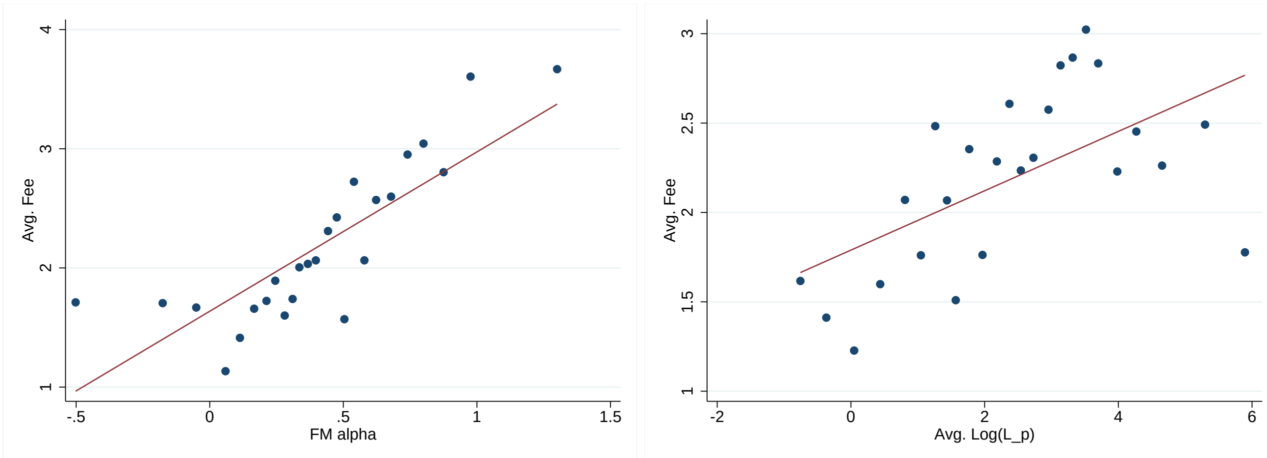

Figure 6 shows bin-scatter plots of average realized fees, calculated as the reported gross rate of return minus the net rate of return. The left-panel plots realized fees against values of $\alpha_i^{p,g}$, the alpha based on the FH+PS model estimated from gross returns and that includes log portfolio illiquidity, $\log(L_p)$, as a characteristic. The right panel plots realized fees against $\overline{\log}(L_p)$, controlling for $\alpha_i^{p,g}$, so that the values of realized fees on the vertical axis are residualized with respect to $\alpha_i^{p,g}$.