The Impact of Hedge Funds on Asset Markets

Mathias S. Kruttli
Department of Economics, University of Oxford and Oxford-Man Institute of Quantitative Finance

Andrew J. Patton
Department of Economics, Duke University, Oxford-Man Institute of Quantitative Finance, and Discipline of Finance at the University of Sydney

Tarun Ramadorai
Saïd Business School, University of Oxford, Oxford-Man Institute of Quantitative Finance, and CEPR

We construct a simple measure of the aggregate illiquidity of hedge fund portfolios, based on the cross-sectional average first-order autocorrelation coefficient of hedge fund returns, and show that it has strong and robust in- and out-of-sample forecasting power for 72 portfolios of international equities, U.S. corporate bonds, and currencies over the 1994 to 2013 period. The forecasting ability of hedge fund illiquidity for asset returns is in most cases greater than, and provides independent information relative to, well-known predictive variables. We rationalize these findings using a simple equilibrium model, in which hedge funds provide liquidity in asset markets. (*JEL* G11, G12, G14, G23)

Introduction

Hedge funds are greatly interesting to a number of different constituencies. Many investors view them as a potential source of low-risk, but outsized, returns. Academics have studied the relationship between hedge funds' returns and systematic risk factors and investigated hedge fund fees, capital accumulation, and disclosure policies. Regulators and policy makers have been wary of the industry ever since the Long-Term Capital Management episode in 1998 that nearly sparked a widespread financial crisis. The periodic

collapses of major hedge funds, such as Amaranth and Madoff, have perpetuated these concerns, leading to perceptions that hedge fund collapses could destabilize underlying asset markets.

These concerns are intuitively plausible: while the global hedge fund industry has around only US$ 2.8 trillion in assets under management (AUM), hedge funds' substantial leverage and the fact that they account for a significant fraction of aggregate trading volume in a range of assets suggests that their impact on markets is large. However, this intuitive plausibility is not backed by a great deal of evidence linking hedge fund activity to underlying asset markets.

Our paper attempts to bridge this gap, providing evidence that hedge funds have a significant and measurable impact on the expected returns of a broad range of assets. We construct a simple measure of hedge funds' aggregate ability to provide liquidity to asset markets and show that time variation in this measure has strong in- and out-of-sample predictive power for a wide range of asset returns spanning three asset classes, namely, international equities, corporate bonds, and currencies. We then build a simple model to shed light on the economic forces that connect hedge fund activity to returns in underlying asset markets. The model helps to explain our empirical results as an equilibrium relationship that naturally arises from hedge funds' provision of liquidity to other market participants.

We focus on liquidity provision as a key hedge fund activity that affects equilibrium in underlying asset markets. This is motivated by much recent work (see, for example, Aragon 2007; Sadka 2010; Aragon and Strahan 2012; and Jylha, Rinne, and Suominen 2014; among others) that shows that hedge funds' return-generation opportunities often arise from liquidity provision to other market participants. One measure of the importance of liquidity-provision activity to hedge funds' enterprise in markets is the ubiquitous use of restrictions, such as lock-up provisions and redemption notice periods on the premature withdrawal of capital by hedge fund investors. These provisions attempt to prevent investors from forcing hedge funds to terminate trades over short horizons when illiquidity-driven price pressure is likely to be most acute; this is a widely-studied phenomenon in the literature on delegated portfolio management (see Shleifer and Vishny (1997) for one early example).

Measuring the illiquidity of hedge funds' holdings is hard, as very limited data about hedge fund positions and turnover are available. To overcome this obstacle, we draw on the widely used model of Getmansky, Lo, and Makarov (2004). They propose that hedge funds holding the most illiquid investments are likely to exhibit the most persistent returns, and provide evidence on the properties of this measure. We therefore measure return autocorrelations using a rolling window of twelve months for each fund in our comprehensive hedge fund database (which tracks the monthly returns of roughly 35,000 hedge funds and spans the period from 1994 to 2013). We then aggregate this measure across funds by cross-sectionally averaging the fund-specific

first-order autocorrelation estimates each month. We find that the resulting index of the aggregate illiquidity of hedge funds, which we dub "ρ", has strong and consistent predictive power for the returns of 72 asset portfolios that span three major asset classes. In-sample, our index is significant for 20 out of 21 international equity indices, 31 out of 42 U.S. corporate bond portfolios spanning the ratings and maturity spectrum, and 6 out of 9 developed-world currencies that we consider.

This predictive power of ρ is not just in-sample; we also find evidence that the measure has appreciable out-of-sample forecasting power. Out-of-sample, our predictor is significantly better than the historical mean return model for 18 out of 21 international equity portfolios, 24 out of 42 U.S. corporate bond portfolios, and 4 out of 9 currencies. In both in-sample and out-of-sample predictive regressions, our proposed measure generally outperforms, and always contributes incremental explanatory power relative to, a range of competitor variables that are known from the relevant asset class-specific literatures to be useful forecasters of asset returns.

We exhaustively eliminate alternative interpretations of the predictive power of ρ that are based on ρ simply proxying for other return predictors known to the literature. Importantly, we verify that ρ does not merely pick up autocorrelation in the returns of the assets that we forecast, or the autocorrelation of a range of factors that are known to be related to hedge fund returns. Our results are also robust to including the lagged asset returns averaged over different windows as competitors. We do find that ρ performs particularly well during the financial crisis. We view this as consistent with our explanation for the source of the predictive power: during this period, capital constraints on hedge fund portfolios ("funding liquidity" à la Gromb and Vayanos 2002 and Brunnermeier and Pedersen 2009) were likely most binding, translating into constraints on hedge funds' liquidity provision to asset markets. We also find in a "placebo" test that our measure of hedge fund illiquidity has no significant predictive power for two extremely liquid assets, namely, government Treasury-bills and 10-year bonds. Finally, momentum in underlying asset returns cannot explain the forecasting performance of ρ, since high values of ρ can be generated by sequences of positive or negative returns, and high values of ρ positively predict returns regardless of the sign of past returns. Overall, our empirical results offer significant support for the importance, and interpretation, of our new hedge fund illiquidity measure.

Our simple model incorporates liquidity constraints into the limits of arbitrage framework of Gromb and Vayanos (2010). In our model, the hedge fund begins with an exogenous endowment of illiquid assets and cash, and generates returns in underlying asset markets by providing liquidity to absorb the random demands of noise traders. The hedge fund thus acts as a market maker for the risky asset. However, the liquidity provision capacity of the hedge fund is limited by the threat of its outside investors withdrawing their funds, thereby forcing it to hold a sufficient quantity of "precautionary" cash

to satisfy potential redemptions. We show that this implies that a hedge fund with a high initial endowment of illiquid assets will be relatively reluctant to purchase the underlying risky asset and relatively eager to sell it. We evaluate the comparative statics in the model as we vary the endowment of illiquid assets held by the hedge fund and find that in equilibrium, this variation is associated with significant variation in the expected returns on the underlying risky asset. The higher the starting illiquidity of the hedge fund's portfolio, the higher are expected underlying asset returns. We also show that the extent to which hedge funds vary their holdings of risky assets in response to changes in the expected cash shortfall maps into the expected $\overline{R}^2$ from our predictive regressions. We calibrate this elasticity of risky asset demand to expected cash shortfalls to be between -1% and -3% from our predictive regressions.

The predictions of the model fit well with simple intuition. For example, we might expect that the level of illiquidity of the hedge fund interacts with the initial level of illiquidity of the underlying assets. When underlying assets are more illiquid for any reason, it seems intuitive that the impact of hedge fund illiquidity will be greater, leading to an even greater ability for ρ to forecast underlying asset returns. Consistent with this intuition, we find that the highest levels of predictability (and hence the greatest hedge fund risky asset demand elasticities to expected cash shortfalls) are found in relatively lower market capitalization and lower turnover international equity markets, high-yield and long maturity U.S. corporate bonds, and currencies with high interest rates and wide bid-ask spreads. As also might be expected, these estimated $\overline{R}^2$ statistics, and hence elasticities, rise considerably during crises.

The literature on hedge funds is fast-growing: many authors, such as Agarwal and Naik (2004), Fung and Hsieh (1997, 2001, 2004), Bollen and Whaley (2009), and Patton and Ramadorai (2013), have documented that hedge funds are significantly exposed to systematic risk, proxied by indices of equity, bond, and options returns. Moreover, Aragon (2007), Sadka (2010), and Cao et al. (2015) help explain hedge fund performance using measures of asset illiquidity and illiquidity risk. However, evidence on the reverse direction– the exposure of markets to hedge funds–has been harder to come by. Two notable exceptions include Kang, Kondor, and Sadka (2014), who provide evidence suggesting that hedge funds affect idiosyncratic risk in equity markets, and Cao et al. (2013), whose work suggests that hedge funds help in the security price formation process, pushing equity returns to be more in line with the efficient frontier. We view our paper as a useful contribution to this emerging literature, providing evidence that is consistent with hedge funds' impact on a variety of asset markets using a relatively easily constructed variable.

Our paper also adds to the literature on the impacts of liquidity and liquidity risk on asset markets by contributing a new empirical variable, ρ, to the pre-existing canon. We view this new variable as a complement to a number of asset-class-specific liquidity measures (see, for example, Acharya and Pedersen 2005; Pastor and Stambaugh 2003; Khandani and Lo 2011; and

Bongaerts, de Jong, and Driessen 2012) that have been utilized extensively in the literature.

1. Data

1.1 Asset return data

Our empirical analysis covers three major asset classes: international equities, corporate bonds, and currencies. Within each asset class, we study a broad range of individual assets. As our hedge fund data are available only monthly, and our total time-series length is under 20 years, we focus on monthly asset return predictability rather than on higher- or lower-frequency returns. For each of the asset classes, we compare the predictive performance of our hedge fund illiquidity measure with that of a range of benchmarks from the literature. Below, we describe this set of "competing" predictor variables utilized for each asset class.

We employ 21 international equity index returns in our analysis, namely, Austria, Australia, Belgium, Canada, Denmark, Finland, France, Germany, Hong Kong, Ireland, Italy, Japan, Netherlands, New Zealand, Norway, Singapore, Spain, Sweden, Switzerland, the United Kingdom, and the United States. We source our data on returns from Kenneth French's Web site. We compute log excess returns on each of these indices over the sample period 1995 to 2013. As "competitor" predictor variables, we employ the dividend yield for each equity index from Kenneth French's Web site, lagged returns, short term rate (3-month government bond yield) from Bloomberg, term spread (difference between 10-year and 3-month government bond yields) computed from Bloomberg and IMF data, and the S&P 500 VIX from Datastream detrended by subtracting the mean over the trailing 12 months (see, for example, Campbell 1991 and Ferson, Sarkissian, and Simin 2003).[1]

Our dataset on U.S. corporate bonds comprises 42 indices from Bank of America–Merrill Lynch, of which 24 are investment-grade portfolios and 18 are high-yield portfolios, with maturities ranging from one year to over fifteen years. Our corporate bond data begin in 1997 and end in 2013.

We compare the return prediction performance of our hedge fund illiquidity measure with that of the default yield spread, lagged returns, short term rate, term spread, and the market-capitalization-weighted (VWM) excess returns on the S&P 500 index from CRSP. In the Internet Appendix we consider alternative competitors proposed by Bongaerts, de Jong, and Driessen (2012).

---

[1] In the Internet Appendix, we consider additional competitors: default yield spread (difference between BAA and AAA-rated corporate bond yields from Moody's), hedge fund flows, S&P 500 VIX, and S&P 500 VIX AR(2) innovations. Goyal and Welch (2008) consider a wider array of predictor variables for the equity premium than those employed here; however, our focus on international equity returns, rather than just U.S. returns, means that the number of available competitors is restricted by data limitations.

189

Our data on currencies comprise nine currency rates against the U.S. dollar, namely, the Australian dollar, Canadian dollar, Euro, Japanese yen, New Zealand dollar, Norwegian krone, Swedish krona, Swiss franc, and British pound. The data are sourced from Bloomberg and cover the period 1995 to 2013 (we use the Deutsche Mark prior to the introduction of the Euro). Our monthly log currency returns are measured as a U.S. dollar per unit of foreign currency.[2] Given the findings of Meese and Rogoff (1983), we expect that these returns will be extremely difficult to predict. The competitor variables for predicting currency returns include the inflation differential[3] (inflation of the respective currency's country minus US inflation) from the OECD database, the one-month Libor interest rate differential (rate of the respective currency minus U.S. rate) from Bloomberg, and lagged currency returns. We consider hedge fund flows as a competitor in the Internet Appendix.

Summary statistics for all three asset classes can be found in Table 1. Given the large number of individual bond return series (42 in total), when reporting summary statistics we group them into "investment-grade" and "high-yield" bonds, as well as separate them by maturity, rather than report results for each bond separately.

1.2 An index of hedge fund illiquidity

To compute our hedge fund illiquidity index, we employ monthly hedge fund returns over the period from January 1994 to December 2013, consolidated from data in the BarclayHedge, CISDM, HFR, Morningstar, and TASS databases. These data represent the most comprehensive set of hedge fund data available from public sources, and comprise a total of 34,861 individual hedge funds, including all births and deaths of funds over the period. The funds in the combined database come from a broad range of vendor-classified strategies, which we consolidate into nine main strategy groups: directional traders, emerging markets, fixed income, funds-of-funds, macro, managed futures, multi-process, relative value, and security selection. Details of the process followed to consolidate these data and detailed mappings between self-reported classifications and these nine broad categories of hedge funds can be found in Patton, Ramadorai, and Streatfield (2015). Reported returns are net of management and incentive fees.

Getmansky, Lo, and Makarov (2004) and Lo (2008) propose using the autocorrelation of a hedge fund's reported returns as a proxy for the illiquidity of its asset holdings. Their rationale is that hedge fund managers attempt to value their portfolios at the end of each month; this method is simple when

---

[2] We also considered excess currency returns (over the interest rate differential) and found that the results were very similar.

[3] For Australia and New Zealand, only quarterly CPI data are available to compute the inflation differential.

Table 1
Summary statistics of asset returns

Panel A: International equities

	Average	SD	Max	Min
Australia	0.650	6.324	15.989	−32.581
Austria	0.465	6.904	17.756	−42.588
Belgium	0.548	5.872	15.331	−36.274
Canada	0.714	5.963	19.901	−31.179
Denmark	0.771	5.859	16.635	−28.925
Finland	0.724	8.955	26.508	−34.564
France	0.568	6.096	13.916	−24.795
Germany	0.516	6.531	19.782	−26.658
Hong Kong	0.508	7.274	27.622	−34.038
Ireland	0.537	6.846	30.572	−26.787
Italy	0.354	7.099	18.914	−27.153
Japan	−0.118	5.336	15.200	−14.452
Netherlands	0.539	6.397	15.452	−34.552
New Zealand	0.339	6.266	14.453	−21.282
Norway	0.660	7.679	17.681	−36.923
Singapore	0.321	7.519	25.642	−33.974
Spain	0.612	6.836	18.231	−26.032
Sweden	0.818	7.483	22.440	−32.318
Switzerland	0.624	4.929	12.130	−16.435
United Kingdom	0.509	4.696	13.208	−22.279
United States	0.589	4.478	10.356	−18.347

Panel B: U.S. corporate bonds

	Average	SD	Max	Min
Rating				
Inv. grade (24 portfolios)	0.344	1.663	15.146	−16.070
High yield (18 Portfolios)	0.458	3.633	31.302	−36.268
Maturity				
1–3Y (7 portfolios)	0.339	1.726	31.302	−22.596
3–5Y (7 portfolios)	0.361	1.843	14.882	−24.184
5–7Y (7 portfolios)	0.375	2.184	20.874	−25.359
7–10Y (7 portfolios)	0.312	2.670	17.898	−30.674
10–15Y (7 portfolios)	0.386	3.115	23.176	−36.268
15 + Y (7 portfolios)	0.565	3.539	24.717	−28.694

Panel C: Currencies

	Average	SD	Max	Min
Australia	0.061	3.615	9.899	−17.108
Canada	0.122	2.401	8.936	−13.009
Euro	0.043	2.949	9.609	−10.196
Japan	−0.024	3.221	16.273	−9.698
New Zealand	0.110	3.636	12.507	−13.931
Norway	0.047	3.161	7.785	−13.739
Sweden	0.063	3.242	9.127	−11.621
Switzerland	0.168	3.149	12.368	−11.945
United Kingdom	0.025	2.409	9.044	−10.215

Reported are the average, standard deviation, maximum, and minimum, for monthly log excess returns in % of international equities and U.S. corporate bonds. For currencies, the statistics are based on monthly returns with US$ as the base currency. For international equities and currencies, the time series starts in January 1995 and ends in December 2013. For U.S. corporate bonds, the time series starts in January 1997 and ends in December 2013.

all assets in the portfolio are liquid and trade often, but managers may be forced to use models to estimate values for illiquid assets in their portfolios. A common approach is to linearly extrapolate prices, leading to reported hedge returns exhibiting positive autocorrelation. Another factor generating

spurious positive autocorrelation is the relationship between the liquid and illiquid assets in hedge fund portfolios, which creates nonsynchroneity of returns (see, for example, Scholes and Williams 1977; Dimson 1979). For these reasons, higher levels of autocorrelation in hedge fund returns are associated with higher levels of illiquidity in hedge fund portfolios.

Getmansky, Lo, and Makarov (2004) also consider the case in which hedge fund managers engage in "performance smoothing;" that is reported returns are a smoothed version of true returns (see Bollen and Pool 2008, 2009 and Agarwal, Daniel, and Naik 2011 for related work). Intentional smoothing leads to the same features as the smoothing that arises from "marking to model;" that is, reported returns will be positively autocorrelated, though in this case, underlying true returns are not autocorrelated. If managers engage in return smoothing, which is unrelated to asset illiquidity, this creates measurement error, which biases against us finding a significant result. If, however, these managers engage in "opportunistic smoothing;" that is, their intentional smoothing is higher when underlying assets are illiquid, providing greater opportunities for misrepresentation, this maintains the direction of the relationship that we expect to see between hedge fund illiquidity and asset returns.[4]

We compute the aggregate measure of hedge fund portfolio illiquidity (which we denote by ρ_t in month t, or simply ρ in what follows) using the following simple procedure. In each month t, for all funds for which there are 12 prior months of return data available, we estimate the first-order autocorrelation of returns. ρ_t is simply the cross-sectional average of these autocorrelations across all funds in each month. In our main analysis we use an equally weighted average of these individual illiquidity estimates, and in our robustness checks we show that using an AUM-weighted measure leads to similar results.

Given the 12-month "burn-in" period for estimating rolling autocorrelations, the time series of ρ begins in December 1994 and ends in December 2013. Figure 1.A plots AUM-weighted and equally weighted ρ over time, and shows that it spikes during the LTCM crisis of 1998 and the "Quant meltdown" of August 2007 (see Khandani and Lo 2011), as well as during the Great Recession. However, the recession following the NASDAQ crash at the turn of the millennium did not appear to affect the measure greatly. It is also evident from the plot that equally weighted ρ has far greater volatility than the AUM-weighted version, which suggests that larger funds may have more room to engage in a greater degree of smoothing behavior. This may not be surprising given that funding liquidity pressure imposed

---

[4] The smoothing mechanisms considered by Getmansky Lo, and Makarov (2004) are related to managerial expectations of one's own performance, without assuming any predictive capability of managers for underlying asset performance in the future. However, if managers do have such predictive ability, and their smoothing behavior is correlated with future underlying asset returns, that might be an additional rationalization of our results. We thank the Editor (Wayne Ferson) for bringing this point to our attention.
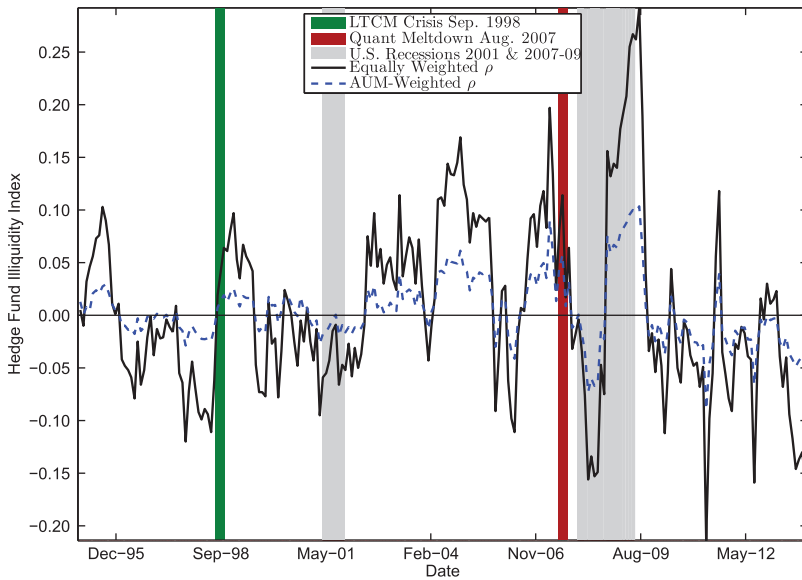
Figure 1 A
Hedge fund illiquidity index over time: This figure shows the AUM-weighted and the equally weighted hedge fund illiquidity indices from January 1995 to December 2013.

by prime-brokerage relationships and capital inflows to funds may be less of an issue for larger, better-established hedge funds, permitting them to manage their reported returns to a greater extent.

Figure 1.B plots ρ and three variations of the Pastor-Stambaugh Liquidity measure. Pastor and Stambaugh (2003) construct a measure of liquidity that is based on the volume-induced return reversals of individual stocks. Their idea is that high volume captures days of high noise trading, à la Campbell, Grossman, and Wang (1993). The essential idea is that the greater the return reversal after a day of high volume, the less liquid is the market; that is, a liquidity provider will receive greater compensation on such days.

The correlation between ρ and the Pastor-Stambaugh measures is -0.06 for the Level of aggregate liquidity, -0.05 for the Innovations in aggregate liquidity, and -0.01 for the Traded liquidity factor. This suggests that our measure captures variation in liquidity arising from a different source than the one identified in the Pastor-Stambaugh measure.

Table 2 reports summary statistics for hedge fund illiquidity indices computed for each individual style, followed by those computed across all funds, regardless of style. The table reveals cross-sectional variation in illiquidity that is consistent with intuition. For example, the fixed income style has a high mean level of illiquidity, while the managed futures funds are the least illiquid.

Panel B of Table 2 presents summary statistics on illiquidity indices formed according to the lock-up or redemption notice period of the fund, for funds
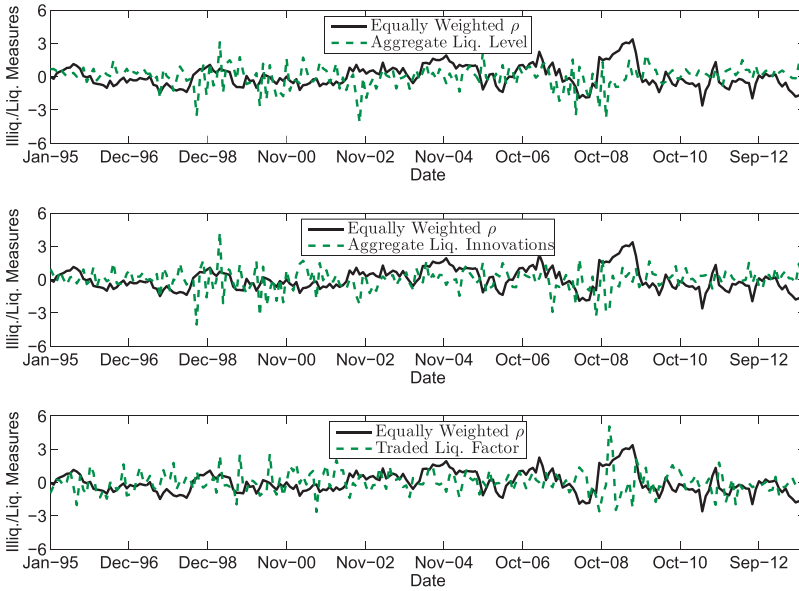
Figure 1 B
Hedge fund illiquidity index and Pastor-Stambaugh Liquidity measures: This figure shows the equally weighted hedge fund illiquidity index from January 1995 to December 2013 plotted against the Pastor-Stambaugh Level of aggregate liquidity, Innovations in aggregate liquidity, and Traded liquidity factor. All variables are de-meaned and divided by their respective standard deviations. The correlation between the equally weighted hedge fund illiquidity index and the Pastor-Stambaugh measures is −0.06 for the Level of aggregate liquidity, −0.05 for the Innovations in aggregate liquidity, and −0.01 for the Traded liquidity factor.

for which this information is available.[5] The table shows that the average level of the equally weighted indices increases in both the lock-up period and the redemption notice period, consistent with our interpretation of this index as measuring illiquidity (see Aragon 2007). The AUM-weighted index levels increase in redemption notice period length, but not lock-up period.

In our analysis, we focus on ρ computed using all funds, rather than using funds belonging to specific styles associated with particular assets. This is partly motivated by a desire for simplicity, and partly by a desire for robustness to potential style drift and misclassification arising from the static and self-reported style classifications in our hedge fund databases. The Internet Appendix shows how self-reported style classifications can create scenarios in which an index based on all available funds delivers better predictive ability than subindices which rely on noisy style classifications. Nevertheless, in Section 4, we show results obtained when employing illiquidity indices based on hedge funds in styles that are, in theory, more closely related to a

---

[5] Both of these characteristics tend to bunch at a few points in the cross-section of hedge funds, and so we divide them into three groups: below median, equal to median, and above median. The "equal to median" group contains 28% of funds for the redemption notice period (median is 30 days) and 51% of funds for the lock-up period (median is 360 days).

Table 2
Summary statistics of hedge fund illiquidity index

| | Panel A: Hedge fund styles | | | | | |
| | Equally weighted | | | AUM-weighted | | |
Styles	Average	SD	Corr. with All	Average	SD	Corr. with All
CTA	−0.073	0.076	0.694	−0.020	0.024	0.649
Directional traders	0.002	0.093	0.849	0.004	0.038	0.829
Emerging markets	0.024	0.116	0.734	0.001	0.049	0.648
Fixed income	0.071	0.075	0.796	0.021	0.033	0.575
Funds of funds	0.059	0.130	0.951	0.023	0.062	0.933
Global macro	−0.055	0.071	0.778	−0.018	0.031	0.639
Multi-strategy	0.037	0.087	0.892	0.017	0.035	0.867
Relative value	−0.041	0.074	0.675	−0.009	0.039	0.402
Security selection	−0.016	0.068	0.878	−0.006	0.027	0.830
All	**0.006**	**0.085**	**1.000**	**0.004**	**0.032**	**1.000**

	Panel B: Hedge funds with lock-up and redemption periods					
Lock-up < median	0.013	0.094	0.846	−0.007	0.045	0.582
Lock-up = median	0.034	0.095	0.931	0.016	0.034	0.881
Lock-up > median	0.044	0.097	0.928	0.010	0.028	0.809
Lock-up all	0.031	0.093	0.941	0.009	0.031	0.904
Redemption < median	−0.022	0.079	0.964	−0.008	0.028	0.911
Redemption = median	0.012	0.078	0.947	0.004	0.027	0.941
Redemption > median	0.057	0.105	0.964	0.021	0.044	0.947
Redemption all	0.021	0.088	0.991	0.008	0.033	0.984
All	**0.006**	**0.085**	**1.000**	**0.004**	**0.032**	**1.000**

For the illiquidity index of each investment style, lock-up period, and redemption period, we report the average, the standard deviation, and the correlation with the All hedge fund illiquidity index. The hedge fund illiquidity index for a specific hedge fund in month t is the autocorrelation of the hedge fund's monthly returns estimated over the last 12 months. The hedge fund illiquidity index shown in the table is an equally or AUM-weighted average across all the hedge funds with a particular investment style, lock-up period, or redemption period. The time series starts in January 1995 and ends in December 2013.

given asset class and find that they yield similar results to using the aggregate measure.

We note here that while Figure 1.A reveals ρ to be reasonably persistent (its first-order autocorrelation is 0.84, with a 95% confidence interval of [0.72,0.96]), an augmented Dickey-Fuller test for a unit root strongly rejects the null in favor of stationarity. Further, the point estimate is sufficiently far from unity so that it does not represent a cause for concern about standard statistical inference. We also note that the characteristics of our data are such that the potential bias arising from our use of a regression with a persistent predictor variable (Equation (1), below) is small. Stambaugh (1999) shows that when the predictor variable (ρ, in our case) is persistent, and the innovations to the predictor variable are (negatively) correlated with the shocks to the target variable (returns in our case), the finite-sample OLS estimates of the predictor variable's coefficient can be upward biased. In our data, the predictor variable has an autocorrelation of 0.84, compared with 0.95 to 0.99 in Stambaugh (1999), and the correlation between shocks is only around −0.15 in our data, compared with −0.90 to −0.94 in Stambaugh's case.

Implementing the bias adjustment proposed by Stambaugh (1999) for our data leads to a change only in the fourth decimal place, at most, confirming the lack of bias here.

2. Results

We now present our main analysis of the predictive power of the hedge fund illiquidity measures for future asset returns. We study this relationship using a variety of simple models. Our first set of results are based on a full-sample estimation of the relationship, and followed by an out-of-sample forecasting analysis.

2.1 Does hedge fund illiquidity predict asset returns?

Our first analysis of the relationship between hedge fund illiquidity and asset returns is based on a simple predictive regression, in which we use the hedge fund illiquidity measure, denoted ρ_t, to predict one-month-ahead returns $r_{i,t+1}$ on asset i:

$$r_{i,t+1} = \alpha_i + \gamma_i \rho_t + \epsilon_{i,t+1}. \tag{1}$$

For comparison, we also estimate the same univariate model using various competitor variables, suggested by the asset-class-specific literatures, in place of ρ_t. We also include hedge fund flows as a competitor in the Internet Appendix, as they affect the liquidity of a hedge fund's portfolio, and Jylha, Rinne, and Suominen (2014) find that increases in hedge fund flows reduce the amount of short-term return reversal and volatility in the U.S. equity market. Moreover, our model (below) discusses the different roles of flows and the starting endowment of a fund's illiquid assets, making hedge fund net capital inflows a natural competitor for ρ. Flows have a negative correlation of -0.2 with ρ, implying that fund inflows improve the liquidity of a hedge fund's portfolio.[6]

The results of estimating this model on international equity indices are presented in panel A of Table 3, where the standard errors of Newey and West (1987) are employed, as in all of our in-sample time-series regressions.[7] This panel shows that the coefficient on ρ is positive and significant at least at the 10% level for 20 of 21 countries, generating an $\overline{R}^2$ of between 0.7% (Germany) and 7.5% (Australia), with an average of 2.7%.

---

[6] We compute these annual flows fund-by-fund using the standard approach of taking differences in end-of-month AUM accounting for growth in assets arising from returns (see, for example, Fung et al. 2008) over a 3-month and 12-month rolling window, winsorize the resultant flows at the 1% and 99% percentile points, and AUM-weight them to construct an aggregate flow index. We find that AUM-weighted flows generally have greater predictive power than equally weighted flows.

[7] We select the number of lags for the Newey-West standard errors based on the rule $T^{1/3}$. In the Internet Appendix, we check the robustness of our results for variations in the number of lags.

Table 3
In-sample predictive performance

Panel A: International equities

Country	Single predictor						Multiple predictors	
	Hedge fund illiq. index	Dividend yield	Lagged returns	Short-term rate	Term spread	VIX detrended	Competitors excl. HF illiq. index	Competitors incl. HF illiq. index
Australia	(+) 7.487**	(+) 2.304**	(+) 0.558	(−) 0.286	(+) 0.269	(−) −0.216	2.410	(+) 8.820**
Austria	(+) 3.544**	(−) −0.331	(+) 5.400*	(−) 3.118*	(+) 2.930	(−) 2.286**	7.120	(+) 9.833**
Belgium	(+) 1.883***	(−) 0.999*	(+) 6.517**	(−) 2.244**	(+) 2.909*	(−) 2.500**	9.555	(+) 11.018**
Canada	(+) 2.969***	(+) −0.021	(+) 1.210	(−) −0.442	(+) −0.251	(−) 0.421	0.189	(+) 2.805**
Denmark	(+) 2.033**	(+) 0.847	(+) 0.864	(−) 1.925**	(+) 2.809**	(−) 1.230*	3.847	(+) 6.160**
Finland	(+) 0.953**	(+) −0.444	(+) 3.234**	(−) 0.900	(+) 2.237**	(+) −0.277	6.768	(+) 7.270*
France	(+) 1.293**	(−) −0.236	(+) 0.546	(−) 0.830*	(+) 1.500*	(−) 0.120	0.923	(+) 2.325**
Germany	(+) 0.691*	(−) −0.022	(+) −0.064	(−) 1.472**	(+) 1.423**	(−) 0.056	1.190	(+) 1.664*
Hong Kong	(+) 4.278**	(+) −0.222	(+) 0.844	(−) −0.383	(+) −0.269	(−) −0.431	−0.603	(+) 2.972**
Ireland	(+) 0.938	(+) −0.420	(+) 3.062	(−) 0.616	(+) 0.515	(−) −0.069	2.425	(+) 3.111*
Italy	(+) 0.874*	(−) −0.082	(+) −0.360	(−) −0.349	(+) −0.325	(−) −0.344	−1.623	(+) −0.790*
Japan	(+) 2.267**	(+) −0.324	(+) 2.383**	(−) 1.109	(−) −0.387	(−) 1.549**	2.211	(+) 4.393**
Netherlands	(+) 2.307**	(+) −0.443	(+) 1.522	(−) 1.285**	(+) 1.648*	(−) 0.613	1.835	(+) 3.850**
New Zealand	(+) 4.089**	(+) −0.388	(−) −0.405	(−) 1.717**	(+) 1.783**	(−) 1.250**	2.113	(+) 7.349**
Norway	(+) 4.062**	(−) 0.072*	(+) 0.810	(−) 2.365**	(+) 2.326**	(−) 0.891	2.556	(+) 5.839**
Singapore	(+) 5.985**	(−) −0.435	(+) 0.890	(−) −0.346	(−) −0.291	(−) −0.036	−0.360	(+) 4.961**
Spain	(+) 2.690**	(+) −0.400	(+) 0.233	(+) −0.444	(+) 0.013	(+) −0.349	−0.011	(+) 2.832**
Sweden	(+) 3.533**	(−) −0.078	(+) 0.984	(−) 0.852	(+) 7.662**	(−) −0.134	7.338	(+) 8.604**
Switzerland	(+) 1.376**	(−) −0.237	(+) 1.334*	(−) 1.564*	(+) 2.403*	(−) 1.021*	2.810	(+) 3.541**
United Kingdom	(+) 2.813**	(+) 1.224*	(+) 2.420	(−) 0.044	(+) 0.576	(−) 0.908	2.734	(+) 5.134**
United States	(+) 1.027**	(+) 1.439**	(+) 0.626	(−) −0.444	(−) −0.406	(−) −0.272	0.762	(+) 1.392*
Across 21 count.	**2.719 (20/0)**	**0.133 (3/2)**	**1.533 (5/0)**	**0.853 (0/9)**	**1.384 (10/0)**	**0.510 (0/6)**	**2.580**	**4.909 (21/0)**

(continued)

Table 3
Continued

Panel B: U.S. corporate bonds

Rating / Maturity (# of portfolios)	Single predictor						Multiple predictors	
	Hedge fund illiq. index	Default yield spread	Lagged returns	Short-term rate	Term spread	VWM excess return	Comp. excl. HF illiq. index	Comp. incl. HF illiq. index
Inv. grade (24)	2.29 (14/0)	2.30 (15/0)	1.93 (13/0)	0.28 (0/6)	1.59 (20/0)	1.36 (0/9)	6.68	7.42 (8/0)
High-yield (18)	5.00 (17/0)	2.58 (7/0)	7.94 (14/0)	1.84 (0/16)	0.78 (5/0)	6.96 (18/0)	12.64	15.34 (18/0)
1–3Y (7)	3.04 (5/0)	4.32 (6/0)	4.26 (5/0)	0.68 (0/4)	0.53 (2/0)	2.66 (3/1)	10.61	11.27 (4/0)
3–5Y (7)	3.73 (5/0)	3.34 (4/0)	6.51 (6/0)	1.53 (0/4)	1.49 (4/0)	4.79 (3/2)	11.37	12.76 (3/0)
5–7Y (7)	3.64 (5/0)	2.43 (3/0)	5.73 (5/0)	1.25 (0/5)	1.52 (6/0)	4.73 (3/1)	10.63	11.81 (4/0)
7–10Y (7)	3.75 (6/0)	1.74 (3/0)	3.99 (6/0)	1.11 (0/4)	1.37 (4/0)	3.51 (3/2)	8.05	10.01 (4/0)
10–15Y (7)	2.57 (4/0)	1.10 (4/0)	3.64 (3/0)	0.47 (0/2)	1.14 (5/0)	3.50 (3/2)	7.94	9.52 (4/0)
15 + Y (7)	3.96 (6/0)	1.57 (2/0)	2.92 (2/0)	0.64 (0/3)	1.42 (4/0)	3.37 (3/1)	7.16	9.47 (7/0)
Across all (42)	**3.45 (31/0)**	**2.42 (22/0)**	**4.51 (27/0)**	**0.95 (0/22)**	**1.24 (25/0)**	**3.76 (18/9)**	**9.23**	**10.82 (26/0)**

Panel C: Currencies

Currencies	Single predictor				Multiple predictors	
	Hedge fund illiq. index	Inflation differential	Interest rate differential	Lagged returns	Competitors excl. HF illiq. index	Competitors incl. HF illiq. index
Australia	(+) 6.310**	(+) 1.443**	(+) 0.094	(+) -0.002	0.608	(+) 5.917**
Canada	(+) 4.052**	(+) 3.269**	(+) -0.218	(−) 0.116	3.587	(+) 5.966**
Euro	(+) 1.158**	(+) 5.497**	(+) 0.160	(+) -0.429	4.406	(+) 5.014
Japan	(+) -0.097	(+) -0.378	(+) -0.329	(+) -0.263	-1.476	(+) -1.212
New Zealand	(+) 3.743**	(+) 3.687**	(−) -0.444	(+) -0.407	2.862	(+) 4.934**
Norway	(+) 1.633**	(+) -0.132	(+) -0.317	(+) -0.417	-1.404	(+) 0.479**
Sweden	(+) 1.362**	(+) 3.293**	(+) 0.180	(+) -0.291	2.212	(+) 3.474**
Switzerland	(+) -0.233	(+) 3.298**	(+) 0.813*	(−) -0.344	3.791	(−) 3.472
United Kingdom	(+) -0.120	(+) 0.132	(−) -0.126	(+) -0.309	0.792	(+) -0.592
Across 9 currencies	**1.979 (6/0)**	**2.233 (6/0)**	**-0.021 (1/0)**	**-0.261 (0/0)**	**1.864**	**3.050 (5/0)**

For international equities and currencies, reported are the signs of the predictors (in the case of multiple predictors, it stands for the sign of the hedge fund illiquidity index), the adjusted R^2 of the in-sample predictive regression in %, and the significance of the hedge fund illiquidity index (* for 10% and ** for 5%). The last row shows the average adjusted R^2, the total number of series for which the coefficient of the predictor is significant (at the 10% level) and positive, and the total number of series for which the coefficient of the predictor is significant and negative (in the case of multiple predictors, only the coefficients of the hedge fund liquidity index are considered). For U.S. corporate bonds, the number of portfolios for which the hedge fund illiquidity index is significant at the 10% level is reported (positive/negative), and the adjusted R^2 is averaged across all portfolios. For international equities and currencies, the time series starts in January 1995 and ends in December 2013. For U.S. corporate bonds, the time series starts in January 1997 and ends in December 2013. Newey-West standard errors are used.

How economically significant are these $\overline{R}^2$ values? We use the model to shed light on this issue. Once we calibrate several parameters of the model to empirical moments, we find that the theoretically predicted $\overline{R}^2$ from the model is highly dependent on λ, which represents the sensitivity of the hedge fund's demand for risky assets to the cash shortfall, that is, the shortfall between a possible redemption request, and the hedge fund's ability to satisfy this redemption request by providing cash to outside investors in response. As we increase λ in the model, the hedge fund's response to the shortfall penalty increases in the sense that the fund will respond to a given shortfall penalty by shedding more of its risky assets.

In reality, this response parameter λ will depend on a range of factors that we do not explicitly consider in our simple model, such as margin requirements faced by the hedge fund, and the perceived illiquidity of the risky asset. However, we can use the model's insight to make some important predictions. For one, the model would suggest that during a crisis, when funding is tight, hedge funds are likely to react more sensitively to a shortfall in cash, implying a greater λ, and consequently a higher $\overline{R}^2$. In support of this prediction, we find that the hedge fund illiquidity index has more predictive power during the financial crisis.

We use the model to attempt to uncover the values of λ that are implied by the $\overline{R}^2$ values that we observe in our empirical exercise. This provides a way to ascertain the economic significance of the $\overline{R}^2$ statistics that we find, in the sense that we can map this back to a deeper microeconomic parameter λ and check whether we obtain theoretically reasonable values for this parameter that are consistent with the observed empirical results. In Table A.2 in the theoretical appendix, we observe that to match the range of $\overline{R}^2$ values that we observe across different assets, the model generates implied values of λ, which map to shortfall elasticities of demand (shown in the last column of Table A.2) between -1% and -3%. In other words, for a 1% change in the cash shortfall, to match the empirically observed $\overline{R}^2$ statistics, we need a percentage change in the expected holdings of the risky asset of between -1% and -3%. This is economically quite large, for the representative hedge fund, since an expected change in holdings of this amount will translate into substantial reductions in trading volume given the high turnover of hedge funds.

There is an alternative way to assess the economic significance of the $\overline{R}^2$ statistics that we uncover. This is the usual Campbell and Thompson (2008) fashion, in which the R^2 maps to a utility gain associated with predictability. These authors show that even with a seemingly tiny out-of-sample (OOS) R^2 of 0.0025, an investor with CRRA utility and a risk aversion coefficient of three experiences a certainty equivalent gain of roughly 1% when predicting the U.S. index. Taking into account that we predict multiple indices across multiple asset classes, and that our OOS R^2 statistics are far higher than this level, as we show below, we believe that the economic significance of our predictive strategy is substantial.

Turning to the other results in the table, not one of the five competing predictor variables is able to match the performance of ρ. The term spread is the best performing, with 5 out of 21 index returns having a coefficient on the term spread that is significant at the 5% level, and another five significant at the 10% level. Using the term spread, the average $\overline{R}^2$ from the model is 1.4%. The lagged returns are the best competitor variable in terms of average $\overline{R}^2$ with 1.5%. The short-term rate performs slightly worse than the term spread with 9 out of 21 index returns being significantly predicted at least at the 10% level and an average $\overline{R}^2$ of 0.9%. The dividend yield and the VIX detrended have five and six significant coefficients (10%) level, respectively. We also consider aggregate hedge fund flows, the VIX level, and the VIX AR(2) innovations as competitors in the Internet Appendix. However, we find that the these predictors perform worse than the predictors shown in Table 3.

Next, we estimate a multivariate regression model in which we include all of the "competitor" predictor variables, together with ρ in a single in-sample predictive regression:

$$r_{i,t+1} = \alpha_i + \beta_i' X_{i,t} + \gamma_i \rho_t + \epsilon_{i,t+1}, \tag{2}$$

where $X_{i,t}$ is a matrix that contains all of the competitor variables. The last two columns of panel A present the results from this model with and without the inclusion of ρ, to reveal additional explanatory power derived from including our measure of hedge fund illiquidity. We see that the average $\overline{R}^2$ jumps from 2.6% to 4.9% with the inclusion of our measure. Furthermore, our measure is positive and significant at least at the 10% level for all 21 country indices, even after the inclusion of the competitor variables. This provides strong support that there is substantial incremental predictive power in hedge fund illiquidity for international equity returns.

The next panel of Table 3 conducts the same in-sample predictive analysis using returns data on U.S. corporate bonds. Given the large number of individual bonds (42) in our sample, we report summaries of the individual regression results, aggregating first by ratings class (investment-grade or high-yield) and then by maturity.[8] For investment-grade bonds, the hedge fund illiquidity index generates an average $\overline{R}^2$ of 2.3% and has a significant coefficient for 14 out of the 24 individual bonds.[9] The default yield spread and the term spread perform better with 15 and 20 significant coefficients, respectively. For high-yield bonds, the hedge fund illiquidity index is significant for

---

[8] Full results for all individual bonds are reported in the Internet Appendix. Additional competitors are also considered in the Internet Appendix.

[9] We investigate whether finding 14 out of 24 individual series to have a significant coefficient is jointly significant when considering the multiple testing aspect of this problem. Taking into account the correlation across individual assets, and using the joint asymptotic distribution across all estimated coefficients, we find that we would expect to see no more than 4 out of 24 with significant coefficients at the 95% confidence level, and so the observed proportion of 14 out of 24 is strongly significant. The proportions we observe for the other asset classes are greater and thus even more significant.

17 out of the 18 portfolios and generates an average $\overline{R}^2$ of 5.0%. The only competitor variable that performs better is the VWM excess return with 18 significant coefficients and an average $\overline{R}^2$ of 7.0%. The lagged returns have a higher average $\overline{R}^2$ with 7.9%, but fewer significant coefficients than the hedge fund illiquidity index (14 compared to 17).

Turning to the results from the multiple regression model, we see that the hedge fund illiquidity index adds substantial predictive power beyond the four competitor variables: the $\overline{R}^2$ rises from 6.7% to 7.4% for investment-grade bonds, and from 12.6% to 15.3% for high-yield bonds. The number of bonds with significant coefficients on the hedge fund illiquidity index is 8 out of 24 for investment-grade bonds, and 18 out of 18 for high-yield bonds. The hedge fund illiquidity index is a significant predictor of U.S. corporate bond returns, and particularly so for high yield bonds.

The lower part of panel B of Table 3 presents results across bond maturity. We find that the results are robust across both the short and long ends of the yield curve, with a slight increase in predictive ability as measured by $\overline{R}^2$ for the longest maturity instruments. We further explore the cross-sectional variation in predictive ability along both the maturity and the ratings dimensions further when we test the predictions of the equilibrium model presented in the next section.

Panel C of Table 3 presents the results of in-sample regressions for currency returns. As is well known (see, for example, Meese and Rogoff 1983 and a large literature thereafter), currencies are generally very hard to predict, and this is confirmed by our analysis. However, our hedge fund illiquidity index performs well: in univariate regressions it generates an average $\overline{R}^2$ of 2.0% (slightly lower than that for the inflation differential, with 2.2%), and is statistically significant for six out of the nine currencies under consideration. In the multiple regression results we see that the $\overline{R}^2$ increases from 1.9% to 3.1% with the inclusion of ρ, which is significant for five out of the nine currencies. Despite the notoriously challenging task of predicting currency returns, it appears that ρ appears to be able to add some value.

Figure 2 shows the length of time over which ρ retains predictive ability for future equity, bond, and currency returns. The top panel of the figure shows the results for the single-variable predictive regressions, and the bottom panel shows the results for the multiple-variable predictive regressions. The line in each figure shows the average $\overline{R}^2$, and the bars show the number of assets for which ρ is statistically significant. Along the horizontal axis, we vary the forecast horizon: $h = 1$ corresponds to our baseline regression, in which we forecast returns using one-month lagged ρ, and $h = 6$ corresponds to forecasting returns using six-month lagged values of ρ. The figure shows that the forecasting power of ρ is quite long-lasting: the forecasting performance for all three asset classes deteriorates significantly at the seven- to eight-month horizons.
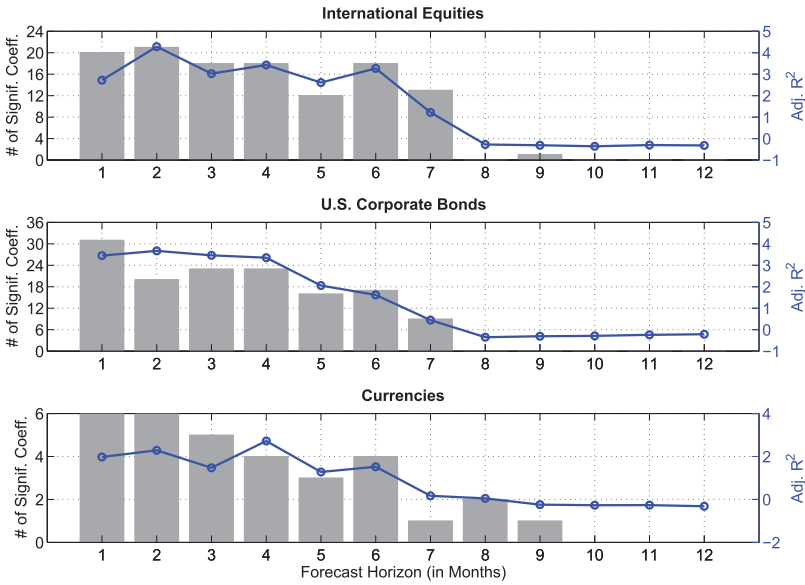
Figure 2 A
Single predictor forecast horizon: This figure shows the average adjusted R^2 (line, right y-axis) and number of significant coefficients (bars, left y-axis) of the hedge fund illiquidity index for different forecast horizons (in months). The results are for single predictor in-sample regressions, that is, the only predictor is the hedge fund illiquidity index.
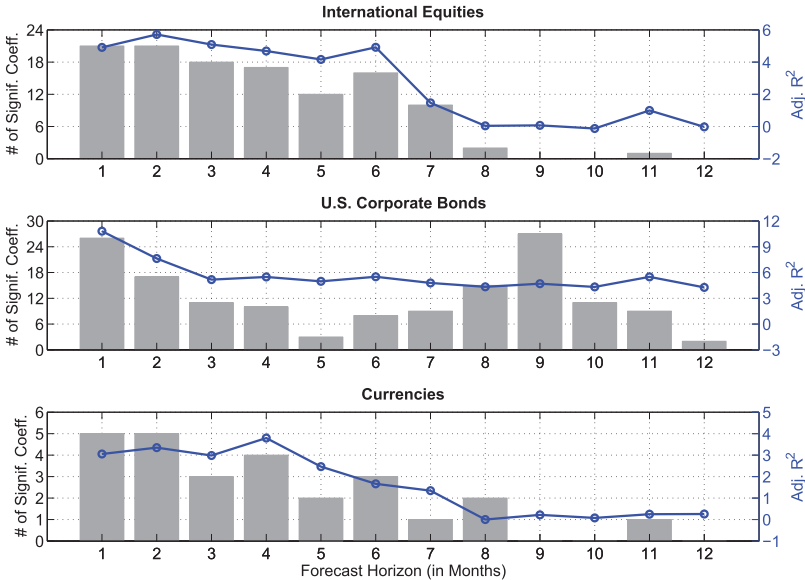
Figure 2 B
Multiple predictors forecast horizon: This figure shows the average adjusted R^2 (line, right y-axis) and number of significant coefficients (bars, left y-axis) of the hedge fund illiquidity index for different forecast horizons (in months). The results are for multiple predictor in-sample regressions, that is, the hedge fund illiquidity index is included together with competitor variables.

This result on the long-lasting nature of the forecasting power is consistent with our explanation for the predictive power of ρ; it could easily take up to six months for an illiquid hedge fund to unwind its portfolio before it can recommence providing liquidity to asset markets. We motivate this insight more rigorously in our discussion of the model in Section 3.

Table 4 reports the number of significantly positive and significantly negative coefficients on ρ in-sample, when we include three important additional control variables. These are the 12-month rolling average asset returns, autocorrelations of the asset returns, and autocorrelations of risk factors, which help to explain hedge fund returns. We include these controls on the right-hand side in order to assuage concerns that ρ simply proxies for attributes of the underlying asset returns or risk factors. The top row of the table (labeled "No controls") reports the number of positive and negative significant coefficients for the single predictor in-sample regressions and is repeated from Table 3 for ease of comparison. Further below, the table shows that the performance of ρ is virtually unchanged by the addition of the average asset returns or the autocorrelation coefficients of the asset returns themselves. The results are similarly robust when including the autocorrelation coefficients of the risk factors,

It is also worth noting here that these results cannot be explained by momentum in underlying asset returns. Time-series momentum corresponds to positive autocorrelation in asset returns, meaning that positive or negative returns will predict returns that rise or fall in the future, respectively. In contrast, high values of ρ are generated by continuations of returns regardless of whether these returns are positive or negative. We find, regardless of the sign of past returns, that high ρ predicts higher future returns.

In panel B of Table 4 we undertake a "placebo" test, where we attempt to use our hedge fund illiquidity index to predict returns on extremely liquid assets. If our rationalization of the predictive capability of ρ is correct, we should not be able to predict the returns of very liquid assets using this measure. We consider excess returns on Treasury bills and 10-year bonds for ten countries (with data taken from Global Financial Data and Bloomberg) for this test and show that ρ has virtually no predictive power for either of these assets in any country, providing support for our interpretation of the mechanism through which ρ predicts asset returns.

2.2 The out-of-sample predictive power of hedge fund illiquidity

The above analyses used the full sample of data to estimate the relationship between ρ and asset returns. In this section we consider the out-of-sample performance of models based on ρ, as well as the performance of competitor predictor variables. Out-of-sample evaluation is increasingly an important metric for predictive models of asset returns, on account of its greater relevance for real-world investors (see, for example, Goyal and Welch 2008, and

Table 4
Other competitors and assets

Panel A: Additional control variables

Control variables	Significance of HF illiq. index (# of pos./neg. coeff.)		
	Int. equities	U.S. corp. bonds	Currencies
No controls	(20/0)	(31/0)	(6/0)
12M average asset return	(20/0)	(34/0)	(6/0)
12M asset return autocorrel.	(20/0)	(31/0)	(6/0)
12M risk factor autocorrel.			
HML U.S.	(20/0)	(30/0)	(6/0)
Mkt-Rf U.S.	(21/0)	(26/0)	(5/0)
Momentum U.S.	(20/0)	(31/0)	(6/0)
SMB U.S.	(20/0)	(31/0)	(6/0)
HML Global	(20/0)	(30/0)	(6/0)
Mkt-Rf Global	(21/0)	(31/0)	(5/0)
SMB Global	(18/0)	(31/0)	(6/0)
WML Global	(20/0)	(31/0)	(6/0)
PTFSBD	(18/0)	(28/0)	(6/0)
PTFSCOM	(21/0)	(27/0)	(4/0)
PTFSFX	(19/0)	(27/0)	(6/0)

Panel B: HF illiquidity index in-sample predictability for government bonds

	Treasury bills	10-year bonds
Australia	(+) −0.314	(−) −0.068
Canada	(+) −0.251	(+) −0.432
Denmark	(+) −0.041	(−) −0.443
Germany	(−) −0.440	(−) −0.444
Japan	(+) −0.062	(−) −0.209
Norway	(−) 0.248	(+) −0.247
Sweden	(−) 0.145	(−) −0.441
Switzerland	(−) −0.173	(+) −0.380
United Kingdom	(+) −0.184	(+) −0.253
United States	(+) −0.418	(+) −0.397
Across 9 countries	**−0.149 (0/0)**	**−0.331 (0/0)**

Panel A of this table presents the significance at the 10% level of the hedge fund illiquidity index when adding the 12-month average asset returns, the 12-month autocorrelation of the asset returns, or the 12-month autocorrelation of various risk factors, as controls to the single variable predictive regression. The first number stands for the number of assets with positive significant coefficients on the hedge fund illiquidity index and the second for negative significant coefficients. The dependent variables are the monthly log excess returns, except for currencies, where we use the monthly log return. The time series for U.S. corporate bonds starts in January 1997 and ends in December 2013. For currencies and international equities, the time series starts in January 1995 and ends in December 2013. Panel B presents the predictive power of the hedge fund illiquidity index for monthly log excess returns on government bonds in a single variable predictive regression. The time series starts in January 1995 and ends in December 2013. For Norway, Sweden, and Switzerland, bonds with 7- to 10-year maturity are used instead of 10-year bonds. The sign of the hedge fund illiquidity index, the adjusted R^2 of the in-sample predictive regression in %, and the significance of the hedge fund illiquidity index (* for 10% and ** for 5%) are reported. The last row shows the average adjusted R^2, the total number of assets for which the coefficient of the predictor is significant (at the 10% level) and positive, and the total number of assets for which the coefficient of the predictor is significant and negative.

Campbell and Thompson 2008 for analyses of the predictive ability of various models for the equity premium). In this subsection we take the perspective of a real-world investor who attempts to forecast asset returns in real time using ρ.

Given the well-known difficulty of out-of-sample asset return prediction, we focus on simple single-variable predictive models, as in Equation (1). We estimate the parameters of the models using a rolling window of 60 months of data. So for international equities and currencies, our out-of-sample period is January 2000 to December 2013. For U.S. corporate bonds, the out-of-sample period is January 2002 to December 2013. We report the OOS R^2 from each of the models, and we modify the Clark and West (2006) test to enable formal comparison of the performance of models employing ρ against a model that only includes a constant. (The original Clark-West test requires the smaller model to contain no parameters at all, whereas for asset returns it is more reasonable to take a model with just a constant term as the benchmark. Clark and West (2007) consider comparisons of general nested models, but do not obtain explicit asymptotic critical values for their test statistic.) The details of this simple extension are presented in Appendix B, and a small simulation study in the Internet Appendix verifies that it performs well in finite samples.

Panel A of Table 5 presents the out-of-sample results for predicting international equity index returns. Forecasts based on ρ generate an average OOS R^2 of 1.9%, ranging from -0.5% for Japan to 6.7% for Australia. The OOS R^2 is negative in only two cases, and for 18 out of the 21 countries, the extended Clark-West test rejects the constant model in favor of a model based on ρ.[10] In contrast, out-of-sample forecasts based on the competitor variables (dividend yield, lagged returns, short-term rate, term spread, and detrended VIX) generate mostly negative average OOS R^2 and are significant for a maximum of 13 countries in the sample. These results provide support for the result obtained in our full-sample analysis that ρ is useful for predicting international equity index returns.

In panel B of Table 5 we find similarly good results for U.S. corporate bonds. The average OOS R^2 from forecasts based on ρ is 2.8% and is positive for all ratings and maturity subsets of these bonds. We also formally reject the constant model in favor of our model for 24 out of the 42 bonds. The forecasts based on the competitors (default yield spread, lagged returns, short-term rate, term spread, and VWM excess return) all perform worse: all five generate negative average OOS R^2. Forecasts based on the term spread perform the best of these five variables in terms of OOS R^2, with an average of -0.6%, and significantly beating the benchmark constant model for 20 out of 42 bond portfolios. For the VWM excess return, the only subset of bonds for which the average OOS R^2 is positive is the high-yield group, and even there the performance of this predictor variable is inferior to ρ.

---

[10] Note that a higher OOS R^2 does not necessarily translate to greater significance in the Clark-West test. For example, the forecasts based on the hedge fund liquidity index have an OOS R^2 of 2.3% for Sweden, and we reject the null that the constant model is best at the 10% level, while for Hong Kong the OOS R^2 is 2.1%, and we reject the null in favor of the larger model at the 5% level.

Table 5

Out-of-sample forecast performance

	Panel A: International equities					
Country	Hedge fund illiq. index	Dividend yield	Lagged returns	Short-term rate	Term spread	VIX detrended
Australia	6.672**	−2.717	−1.650	−1.881	−2.585	−4.861
Austria	2.732**	−1.196	0.972	6.076*	4.229	−3.899
Belgium	2.094*	−1.500	4.437*	7.596**	2.858*	−2.072
Canada	2.093**	−3.352	−3.519	−3.964	−3.390	−4.282
Denmark	1.481*	−2.348	−0.658	2.247*	0.865*	−2.860
Finland	0.692**	−5.907	1.694**	0.529**	−2.263*	−1.421
France	0.784*	−7.682	−1.637	3.965**	−0.073	−4.029
Germany	−0.154	−15.602*	−2.456	1.744**	0.123	−3.514
Hong Kong	2.082**	−4.849	−1.194	−5.918	−7.908	−11.102
Ireland	0.088*	−7.887	2.089	8.286**	−1.337	−1.172
Italy	0.058*	−3.154	−2.038	2.393*	−0.194	−3.718
Japan	−0.465	−2.396	−0.160	−4.607	−4.197	−2.510*
Netherlands	2.040*	−8.383	−2.656	2.373**	−1.513	−5.435
New Zealand	4.629**	−41.890	−2.490	1.257	−0.328	−0.346
Norway	3.524**	−4.288	−3.290	5.697**	6.105**	−1.955
Singapore	2.892**	−6.545	−2.715	−4.024	−3.076	−8.605
Spain	1.954**	−4.150	−2.695	1.557*	−0.046	−4.673
Sweden	2.269*	−10.378	−1.463	5.341**	5.748**	−5.016
Switzerland	0.779	−5.740	−0.046	4.149*	−0.438	−0.770
United Kingdom	2.364*	−6.655	−0.003	−2.675	−6.210	−1.176
United States	1.021*	−0.980	−1.317	−5.304	−4.931	−0.888*
21 countries	**1.887 (18)**	**−7.029 (1)**	**−0.990 (2)**	**1.183 (13)**	**−0.853 (5)**	**−3.540 (2)**

	Panel B: U.S. corporate bonds					
Rating/Maturity	Hedge fund illiq. index	Default yield spread	Lagged returns	Short-term rate	Term spread	VWM excess return
Inv. grade (24 port.)	1.853 (7)	−13.592 (0)	−6.170 (1)	−2.916 (4)	0.986 (20)	−6.809 (0)
High-yield (18 port.)	4.029 (17)	−20.749 (0)	−4.156 (10)	−2.412 (1)	−2.704 (0)	2.265 (16)
1–3Y (7 port.)	2.809 (4)	−18.018 (0)	−16.762 (1)	−3.602 (1)	−2.389 (1)	−3.980 (2)
3–5Y (7 port.)	2.900 (3)	−14.167 (0)	−2.141 (3)	−2.497 (1)	−0.389 (4)	−1.909 (3)
5–7Y (7 port.)	2.724 (3)	−15.010 (0)	−3.352 (2)	−2.233 (1)	−0.070 (4)	−2.678 (3)
7–10Y (7 port.)	3.024 (5)	−17.200 (0)	−4.819 (1)	−2.385 (1)	0.109 (4)	−3.728 (2)
10–15Y (7 port.)	2.127 (4)	−16.496 (0)	−0.044 (1)	−3.162 (0)	−0.778 (4)	−2.599 (3)
15+Y (7 port.)	3.132 (5)	−19.064 (0)	−4.721 (3)	−2.321 (1)	0.395 (3)	−2.628 (3)
42 portfolios	**2.786 (24)**	**−16.659 (0)**	**−5.306 (11)**	**−2.700 (5)**	**−0.595 (20)**	**−2.920 (16)**

	Panel C: Currencies			
Currency	Hedge fund illiq. index	Inflation differential	Interest rate differential	Lagged returns
Australia	4.149**	−1.524	−3.421	−2.656
Canada	1.323**	0.759**	−1.835	−1.800
Euro	−1.287	4.742**	−1.441	−2.312
Japan	−2.983	−2.041	−4.620	−3.213
New Zealand	2.567**	−1.954	3.478**	−1.629
Norway	0.060*	−2.906	−1.915	−4.708
Sweden	−1.698	−0.281	−0.909	−3.180
Switzerland	−2.208	−0.432*	−2.269	−1.802
United Kingdom	−2.889	−3.436	0.566	−1.695
9 currencies	**−0.329 (4)**	**−0.786 (3)**	**−1.374 (1)**	**−2.555 (0)**

The results of rolling 5-year OOS forecasts with single predictor regressions are reported. The results for the hedge fund illiquidity index and the competitor variables are shown. The dependent variable is the monthly log excess return for international equities and U.S. corporate bonds, and the monthly log return for currencies. For international equities and currencies, we report the OOS R^2 and whether we can reject the historical average model based on an extension of the Clark-West test-statistic, which is denoted by * for 10% significance and ** for 5% significance. For U.S. corporate bonds, the number of portfolios for which we can reject the historical average model at the 10% level is in parentheses. The last row in every panel shows the average OOS R^2 and the total number of portfolios for which the historical average model is rejected. For international equities and currencies, the time series starts in January 1995 and ends in December 2013. The time series for U.S. corporate bonds starts in January 1997 and ends in December 2013.

Panel C of Table 5 presents results for out-of-sample currency return prediction. As might be expected given the well-documented difficulty of predicting currency returns, the OOS R^2 is lower for this asset class, and indeed for all four variables we find a negative average OOS R^2, with the highest (-0.3%) coming from ρ. For four of these currencies, we are able to reject the constant model in favor of using ρ: Australia, Canada, New Zealand, and Norway. Notably, these are all relatively high interest-rate commodity currencies, which are generally associated with the long side of the frequently studied carry trade (see Burnside, Eichenbaum, and Rebelo 2011 for an overview of the carry trade literature). We explore this finding further when we test our model predictions below. Finally, when we compare the performance of ρ to forecasts of currency returns based on competing variables, the primary significant result is the usefulness of the inflation differential for predicting the euro return: this forecast generated an OOS R^2 of 4.7% and is significantly better than the constant model; the interest rate differential rejects the constant model for New Zealand.[11] Lagged currency returns fail to generate a positive OOS R^2 or reject the constant model for all currencies.

It is worth noting that our out-of-sample period spans the financial crisis and the "great recession." Unfortunately, hedge fund data only became available in the early to mid-1990s, and so we are unable to conduct formal analysis of the impact of recessions on the predictive relationships we document here. To gain some insight into the impact, however, we construct a time series of the cumulative sum of squared errors of the historical mean return model and the predictive regression, similar to that used in Goyal and Welch (2008). Plots of these time series are shown in Figure 3. For all three asset classes, we find that ρ performs particularly well during the financial crisis. We view this as consistent with our explanation for the source of the predictive power: during this period, capital constraints on hedge fund portfolios ("funding liquidity" à la Gromb and Vayanos (2002) and Brunnermeier and Pedersen (2009)) were likely most binding, translating into constraints on hedge funds' liquidity provision to asset markets.

The next section presents a simple model that provides a more rigorous foundation for this and our other results, and generates additional empirical predictions, which we subsequently test.

3. Simple Model and Predictions

Consider an economy with three types of assets: cash, a risky asset, and an illiquid asset, which can be traded only over longer intervals. The first two of

---

[11] The inflation differentials for Australia and New Zealand have a look-ahead bias, as we have only quarterly CPI data available.

International Equities

U.S. Corporate Bonds

Currencies
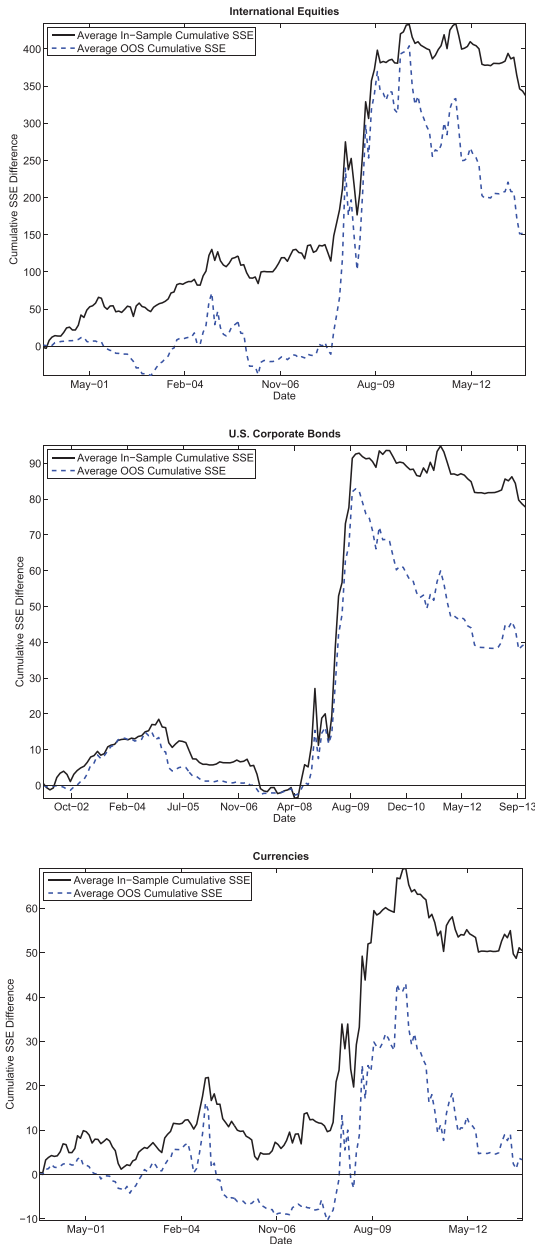
Figure 3
Cumulative forecasting errors: The in-sample line shows the difference of the cumulative sum of squared residuals of a model with just a constant and the residuals of a single variable model with the equally weighted hedge fund illiquidity index as predictor. The OOS line shows the difference of the cumulative OOS squared forecasting errors of the historical average model and the single variable model with the equally weighted hedge fund illiquidity index as predictor. Five-year rolling windows are used. The cumulative differences are averaged across all assets.

these assets, cash and the risky asset, are fungible with one another, but the third is not fungible with either of the first two over short periods of time.

The agents in this economy are a hedge fund, which acts as the market maker for the risky asset, noise traders who subject the asset to demand shocks, and outside investors of the hedge fund, who occasionally demand cash from the fund when they decide to redeem their investments. In this economy, the hedge fund receives compensation for providing liquidity to the noise traders, in the form of return reversals (see, for example, Campbell, Grossman, and Wang 1993).

In this economy, the mechanism that rationalizes our empirical results is the precautionary motive of the hedge fund: the fund can only satisfy redemption requests from outside investors using cash. As a result of the fund's need for cash, the greater the proportion of the highly illiquid asset in the hedge fund's portfolio, the more willing the hedge fund will be to sell the risky asset to the noise traders for cash, and the less inclined the fund will be to buy the risky asset and expend cash.[12] In other words, the (exogenously given) endowment of the highly illiquid asset in the hedge fund's portfolio endogenously generates illiquidity in the risky asset. In Appendix A, we formalize this insight using a model based on the limits of arbitrage framework laid out in Gromb and Vayanos (2010).

This simple mechanism generates four empirical predictions. First, high measured illiquidity of the hedge fund's portfolio (in our empirical implementation, ρ) will positively predict future returns. Second, this predictability for returns will be positive regardless of the direction (buy or sell) of the noise trader demand shock. Third, return reversals will be smaller in magnitude when noise traders buy and larger when noise traders sell, reflecting the asymmetry in the hedge fund's willingness to provide market-making services: the hedge fund is more inclined to sell the risky asset than it is to buy the risky asset. Finally, this mechanism suggests that the above predictions will be even more pronounced when the risky asset is more illiquid, even if this illiquidity arises from unmodeled sources.

In the previous section, we found evidence to support the first of these predictions, namely, that ρ positively predicts future returns on a range of assets. We therefore conduct tests of the remaining predictions of the model. First, we test whether the coefficient on ρ is positive regardless of the sign of the noise trader shock. To identify the sign of these noise trader demand shocks, we adopt the approach of Pastor and Stambaugh (2003), who identify positive and negative noise trader demand shocks using the sign of lagged returns on the underlying asset. We estimate a specification that includes the

---

[12] Note that Ben-David, Franzoni, and Moussawi (2012) find empirically that hedge fund redemption concerns lead to the selling off of more liquid assets. Mitchell, Pedersen, and Pulvino (2007) also find evidence that capital outflows of arbitrageurs can lead to depressed prices that later rebound.

set of all competitor variables X_i for the asset returns in a given class, as well as a dummy for the sign of the lagged return that we interact with ρ:

$$r_{i,t+1} = \alpha_i + \beta' X_{i,t} + \gamma^{\text{NT Sell}} \rho_t I_{\{r_{i,t}<0\}} + \gamma^{\text{NT Buy}} \rho_t I_{\{r_{i,t}>0\}} + \epsilon_{i,t+1}. \quad (3)$$

The prediction of our model is that both $\gamma^{\text{NT Sell}}$ and $\gamma^{\text{NT Buy}}$ are positive. To increase the precision of the parameter estimates, we estimate the above specification as a panel, grouping all assets i within each of the three broad asset classes. We allow for an asset-specific fixed effect in estimation, and we impose that the slope coefficients are the same across all assets within each class. In all panel regressions, standard errors are adjusted for contemporaneous correlation using a delete-cross-section jackknife method (see Shao and Wu 1989 and Shao 1989).

Table 6 presents the results from this model for the three asset classes. Note that Equation (3) nests our specification from Section 2.1, as when $\gamma^{\text{NT Sell}} = \gamma^{\text{NT Buy}}$ the indicator variable drops out. We report results for this special case in the first column of Table 6 for comparison. In the more general model we see that the coefficients $\gamma^{\text{NT Sell}}$ and $\gamma^{\text{NT Buy}}$ are positive for all three asset classes. The estimate of $\gamma^{\text{NT Sell}}$ is always significant, while the estimate of $\gamma^{\text{NT Buy}}$ is not significant for international equities and currencies.

The next prediction of the model is that $\gamma^{\text{NT Sell}}$ should be larger than $\gamma^{\text{NT Buy}}$, reflecting that selling the risky asset to noise traders will alleviate the cash constraints due to possible redemption requests but buying the risky asset has the opposite effect.[13] Table 6 shows that the point estimates from estimating Equation (3) support this prediction for international equities and currencies, but not for corporate bonds. In no asset class, however, is the difference significantly greater than zero. It is worth noting that our use of the sign of the lagged return as a proxy for noise trader demands shocks inevitably introduces additional variability into the model, reducing the power of these tests.

Finally, we test whether ρ has greater predictability for assets that are more illiquid. To test this prediction, we run another fixed effect panel regression for each asset class, in which we again include the set of all competitor variables X_i for the asset returns in a given class, ρ, and ρ interacted with an asset-specific dummy for assets within each broad group that we expect will be more illiquid. As before, the model is estimated for each asset class. The specification is:

$$r_{i,t+1} = \alpha_i + \beta' X_{i,t} + \gamma \rho_t + \gamma^{\text{Illiq.}} \rho_t I_i^{\text{Illiq.}} + \epsilon_{i,t+1}, \quad (4)$$

To assign the $I_i^{\text{Illiq.}}$ dummy, which takes the value of one if portfolio i belongs to a more illiquid subgroup of the asset class, we borrow insight from the

---

[13] Equation (A.1) in the Appendix accounts for this in our model. The equation shows that the shortfall penalty is less likely to be triggered when noise traders buy than when noise traders sell.

Table 6

Panel estimation conditioning on sign of lagged return

Panel A: International equities	(1)	(2)
Dividend yield	−0.088	−0.046
	(−0.370)	(−0.197)
Lagged returns	0.645	0.654
	(1.439)	(1.455)
Short-term rate	−0.277	−0.239
	(−0.861)	(−0.735)
Term spread	0.473**	0.462**
	(2.055)	(2.029)
VIX detrended	−0.059	−0.174
	(−0.114)	(−0.325)
HF illiq. index	1.044**	
	(2.999)	
HF illiq. index × Negative lag ret. dummy		1.057**
		(2.925)
HF Illiq. Index × Positive lag ret. dummy		0.516
		(1.606)
Adj. R^2 (in %)	**5.103**	**5.699**

Panel B: U.S. corporate bonds	(1)	(2)
Default yield spread	0.127	0.130
	(0.485)	(0.502)
Lagged returns	0.390**	0.402**
	(2.178)	(2.064)
Short-term rate	0.020	0.028
	(0.096)	(0.136)
Term spread	0.214	0.221
	(1.189)	(1.238)
VWM excess returns	0.263**	0.268**
	(1.998)	(1.994)
HF illiq. index	0.397**	
	(3.087)	
HF illiq. index × Negative lag ret. dummy		0.267*
		(1.717)
HF illiq. index × Positive lag ret. dummy		0.296**
		(2.813)
Adj. R^2 (in %)	**6.993**	**7.022**

Panel C: Currencies	(1)	(2)
Inflation differential	0.366**	0.365**
	(2.768)	(2.772)
Interest rate differential	0.042	0.042
	(0.301)	(0.298)
Lagged returns	−0.014	0.019
	(−0.081)	(0.108)
HF Illiq. Index	0.351**	
	(2.155)	
HF illiq. index × Negative lag ret. dummy		0.403**
		(2.588)
HF illiq. index × Positive lag ret. dummy		0.126
		(0.883)
Adj. R^2 (in %)	**3.192**	**3.710**

The results of a fixed effects panel estimation with jackknife standard errors clustered by time are reported. The dependent variables are monthly log excess returns or monthly log returns in the case of currencies. The independent variables are lagged by one month and divided by their respective standard deviations. For international equities and currencies, the time series starts in January 1995 and ends in December 2013. The time series for U.S. corporate bonds starts in January 1997 and ends in December 2013. Estimates significant at the 10% level are denoted by *, and estimates significant at the 5% level are denoted by **.

extensive work on illiquidity in a range of asset classes (see, for example, Pastor and Stambaugh 2003; Sadka 2010; and Bongaerts, de Jong and Driessen 2012). For international equities, we use market capitalization and the turnover average over the whole sample period as illiquidity proxies (for values below the median across all equity markets in the sample, the dummy variable takes the value one). For U.S. corporate bonds, high-yield and longer maturity (over 5 years) bonds are considered more illiquid. In the case of currencies, we use the Libor interest rates[14] and the percentage bid-ask spread averaged over the whole sample period as indicators for illiquidity (for values above the median across all currencies in the sample, the dummy variable takes the value one). We expect the coefficient of the illiquidity dummy to be significant and positive, and Table 7 shows that this is indeed the case: $\gamma^{\text{Illiq.}}$ is positive and strongly statistically significant for both measures of illiquidity in each of the three asset classes.

Taken together, this section shows that a simple mechanism in which a market-making hedge fund is sensitive to the threat of redemption requests from outside investors generates empirical predictions that are largely borne out by the data.

The next section outlines a number of robustness checks that we conduct on our results.

4. Robustness

This section presents a variety of robustness checks of our main analyses. The top row of Table 8 (labeled "Base case multiple predictors") corresponds to the bottom-right elements of panels A, B, and C of Table 3. The figures reported are the average $\overline{R}^2$ from the multiple predictor regression, and the number of coefficients (across the individual assets in a given asset class) that are significantly positive and significantly negative at the 10% level.

First, we use AUM-weighting rather than equal weighting, to construct the index. Constructing the index using AUM-weights slightly worsens performance relative to the baseline case, consistent with Figure 1.A, which showed that the AUM-weighted index had smaller fluctuations than the equally weighted index. As noted above, this is not too surprising, as larger funds may be less sensitive to funding liquidity pressures imposed by prime-brokerage relationships and capital inflows.

Second, we check the performance of ρ constructed using funds in style categories that are "close" to the asset class under consideration.[15] For international equities, we consider two hedge fund styles, "Directional traders"

---

[14] We expect that these popular long carry trade currencies will be the first to be avoided in the event of a "flight to liquidity" or to a "safe haven," resulting in their liquidity being lower. See, for example, Campbell, Serfaty-de Medeiros, and Viceira (2010).

[15] Details on the mapping from individual fund styles, as reported to a given hedge fund database, to the ten style categories that we consider in this paper, are provided in Patton, Ramadorai, and Streatfield (2015).

Table 7

Panel estimation with illiquidity dummy variables

Panel A: International equities	(1)	(2)	(3)
Dividend yield	−0.088	−0.087	−0.087
	(−0.370)	(−0.366)	(−0.369)
Lagged returns	0.645	0.644	0.641
	(1.439)	(1.437)	(1.429)
Short-term rate	−0.277	−0.278	−0.276
	(−0.861)	(−0.863)	(−0.857)
Term spread	0.473**	0.474**	0.480**
	(2.055)	(2.060)	(2.079)
VIX detrended	−0.059	−0.059	−0.060
	(−0.114)	(−0.114)	(−0.116)
HF illiq. index	1.044**	0.918**	0.912**
	(2.999)	(2.809)	(2.622)
HF illiq. index × Low mkt. cap. dummy		0.265**	
		(2.312)	
HF illiq. index × Low turnover dummy			0.277**
			(2.497)
Adj. R^2 (in %)	**5.103**	**5.124**	**5.128**

Panel B: U.S. corporate bonds	(1)	(2)	(3)
Default yield spread	0.127	0.126	0.127
	(0.485)	(0.485)	(0.485)
Lagged returns	0.390**	0.392**	0.390**
	(2.178)	(2.189)	(2.174)
Short-term rate	0.020	0.018	0.020
	(0.096)	(0.087)	(0.094)
Term spread	0.214	0.213	0.214
	(1.189)	(1.184)	(1.188)
VWM excess returns	0.263**	0.262**	0.263**
	(1.998)	(1.987)	(2.000)
HF illiq. index	0.397**	0.151	0.265**
	(3.087)	(1.243)	(2.539)
HF illiq. index × High-yield dummy		0.577**	
		(2.468)	
HF illiq. index × Maturity > 5Y dummy			0.199**
			(2.754)
Adj. R^2 (in %)	**6.993**	**7.948**	**7.087**

Panel C: Currencies	(1)	(2)	(3)
Inflation differential	0.366**	0.365**	0.368**
	(2.768)	(2.762)	(2.780)
Interest rate differential	0.042	0.042	0.047
	(0.301)	(0.298)	(0.332)
Lagged returns	−0.014	−0.014	−0.018
	(−0.081)	(−0.082)	(−0.105)
HF illiq. index	0.351**	0.248	0.179
	(2.155)	(1.575)	(1.281)
HF illiq. index × High int. rate dummy		0.234**	
		(2.234)	
HF illiq. index × High avg. spread in %			0.388**
			(3.314)
Adj. R^2 (in %)	**3.192**	**3.283**	**3.524**

The coefficient estimates and t-statistics based on jackknife standard errors clustered by time are reported. The dependent variables are monthly log excess returns, except for currencies, where we use monthly log return. The regressors are lagged by one month and standardized. For international equities, the illiquidity dummy depends on either the total sample average market cap or average turnover standardized by market cap (the asset class median is the benchmark). For U.S. corporate bonds, we use a high-yield dummy and a long-maturity dummy. For currencies, the illiquidity dummy depends on either the total sample average 1M Libor rate or the average percentage spread (the asset class median is the benchmark). The time series starts in January 1995 and ends in December 2013, except for U.S. corporate bonds the start is in January 1997. 10% significance is denoted by * and 5% is denoted by **.

213

Table 8
Robustness checks

Variations of the hedge fund illiquidity index

	Int. equities	U.S. corp. bonds	Currencies
Base case multiple predictors	4.909 (21/0)	10.816 (26/0)	3.050 (5/0)
AUM-weighted	4.891 (21/0)	10.628 (25/0)	2.758 (4/0)
Different styles			
Directional traders	4.150 (19/0)		
Security selection	4.477 (16/0)		
Fixed income		11.055 (33/0)	
Global macro			2.079 (4/0)
Backfill bias robustness			
Delete first 12 months	5.162 (19/0)	10.828 (26/0)	3.247 (5/0)
Delete first 24 months	5.364 (16/0)	10.887 (28/0)	3.176 (5/0)
Offshore/Onshore funds			
Offshore funds	5.113 (21/0)	11.034 (29/0)	2.996 (5/0)
Onshore funds	4.861 (20/0)	10.499 (24/0)	2.983 (5/0)
Size adjustment			
Delete smallest 5% of funds	5.089 (21/0)	11.046 (30/0)	3.023 (5/0)
Delete smallest 10% of funds	4.972 (21/0)	11.035 (29/0)	3.028 (5/0)
Delete smallest 25% of funds	5.012 (21/0)	11.053 (29/0)	3.022 (5/0)
Trimmed	4.570 (21/0)	10.825 (26/0)	2.786 (5/0)
Varying autocorrelation			
AR(2)	2.860 (6/0)	9.918 (21/0)	1.958 (1/0)
MA(1)	5.139 (20/0)	10.562 (24/0)	2.984 (5/0)
MA(2)	3.408 (16/0)	8.083 (26/0)	2.042 (1/0)
Varying window length			
9 months	3.134 (9/0)	10.608 (31/0)	1.855 (1/0)
18 months	3.523 (9/0)	9.726 (19/0)	2.565 (3/0)
24 months	2.916 (3/0)	9.174 (10/0)	2.096 (2/0)

This table presents robustness checks of our main results. The first row corresponds to the bottom-right elements of panels A, B, and C, of Table 3. We present the average adjusted R^2 of multiple predictor regressions across all assets within an asset class and the number of assets for which the coefficient of the hedge fund illiquidity index is significantly (at the 10% level) positive and significantly negative. All predictors are lagged by one month. The dependent variables are the monthly log excess returns, except for currencies, where we use the monthly log return. The time series for U.S. corporate bonds starts in January 1997 and ends in December 2013. For currencies and international equities, the time series starts in January 1995 and ends in December 2013. The earliest starting point for the hedge fund illiquidity index is in January 1995. The AUM-weighted index averages the autocorrelation across hedge fund weighted by AUM. The style-specific hedge fund illiquidity indices are computed with only funds of the corresponding investment style. To control for a potential backfill bias the first 12-months (24-months) of each hedge fund return series are deleted. The offshore (onshore) index is computed with offshore (onshore) hedge funds only. Under size adjustment, the smallest hedge funds are deleted for each month. The trimmed hedge fund illiquidity index sets negative autocorrelations of hedge funds to zero when computing the equally weighted average. Under varying autocorrelation, different measures of autocorrelation are used. Under varying window length, we vary the rolling window over which the autocorrelation of returns for each hedge fund is computed.

and "Security selection." For corporate bonds, the natural style to consider is "Fixed income," and for currencies, we use "Global macro." For international equities, we find that the two style-specific illiquidity indices generate slightly lower numbers of significant coefficients and average R^2. For

currencies, the global macro index also performs slightly worse than the aggregate. For U.S. corporate bonds the fixed income hedge fund illiquidity index generates slightly stronger results than those based on all funds. As mentioned earlier, in the Internet Appendix we present a theoretical analysis of the choice between asset-specific indices and a broader "all" index, highlighting the trade-offs between averaging across more funds and gaining robustness to style misclassification, against greater precision of the illiquidity information.

Third, we test whether our results are robust when we account for a potential backfill bias. We delete the first 12 months and the first 24 months of return observations for each hedge fund and see that the results are very similar to the base case.

Fourth, we construct the hedge fund illiquidity index based on offshore and onshore funds, respectively. The results show that the offshore fund illiquidity index does better for all three asset classes. This result is in line with Aragon, Liang, and Park (2014), who find that offshore hedge funds tend to invest in more illiquid assets, and thus, illiquidity induced trading is more likely for offshore than onshore funds.

Fifth, we delete the smallest 5%, 10%, and 25% of hedge funds, respectively, when constructing the hedge fund illiquidity index. The motivation is to account for the possibility that some of the smaller hedge funds are not viable businesses. We find that deleting the smaller hedge funds from our sample leads to an improved performance for international equities and U.S. corporate bonds. For currencies, the performance is unchanged compared to the base case.

Sixth, we estimate autocorrelations that are "trimmed" to impose nonnegativity, in keeping with the interpretation of ρ as a measure of positive autocorrelation. We see that the index based on trimmed autocorrelations performs slightly worse than our baseline index.

Next, we consider alternative models for obtaining an estimate of the degree of autocorrelation in hedge fund returns. Our baseline model uses simple first-order autocorrelation, which can be interpreted as the autoregressive parameter in an AR(1) model. We consider variations based on an AR(2), MA(1), and MA(2), all estimated using a rolling 12-month window. For the second-order models, there are two parameters that capture autocorrelation, and to summarize these into a single number, we use the R^2 from the model. For the MA(1) model, we simply use the estimated MA parameter. Table 8 shows that the second-order models tend to do worse than the first-order models, likely due to the increased estimation error from the additional parameter. The MA(1) model performs almost as good as our baseline AR(1) model.

Finally, we consider varying the length of the window over which we compute return autocorrelations. Our baseline analysis uses 12 months, and this table considers the use of 9, 18, and 24 months to compute ρ. The trade-off

here is estimation error (longer windows have less estimation error) against timeliness (shorter windows are less "stale"). For international equities and currencies, we see that a 12-month window strikes the optimum balance between these competing goals, outperforming the other choices. The results for U.S. corporate bonds, on the other hand, are less sensitive to the choice of window length, with the window of 9 months outperforming the base case window of 12 months.

5. Conclusion

Detecting evidence of hedge funds' impact on asset markets is an important endeavor given their size, leverage, and significant role in the provision of liquidity. Motivated by the work of Getmansky, Lo, and Makarov (2004), we create a simple time-varying measure of aggregate hedge fund illiquidity, which we dub ρ, by cross-sectionally averaging fund-specific return autocorrelations computed using rolling windows across a large universe of hedge funds. We find that ρ is a highly significant and robust predictor of returns, both in-sample and out-of-sample, for international equity indices, U.S. corporate bonds, and currencies.

We build a simple model of liquidity provision by hedge funds who are endowed with illiquid asset holdings, and face a shortfall penalty for not holding sufficient cash to cover the threat of withdrawals by their outside investors. The model is able to explain our main empirical findings as naturally arising in equilibrium from hedge funds' liquidity provision activity in financial markets and yields additional testable implications which are supported in the data.

We view these results as a useful addition to the literature on hedge funds and the effects of liquidity on asset returns. In future work, we hope to explore the implications of this and other measures of the impact of hedge funds on asset markets to explain a broader range of asset market outcomes.

Appendix A: A Model of Hedge Fund Illiquidity and Asset Return Determination

The model contains three periods, with a time line as follows. In period 0, the hedge fund in our model inherits an illiquid endowment θ, which we assume it cannot sell, and a cash endowment C_0. For simplicity, the return on both the illiquid endowment and the cash endowment are set to zero. In period 1, the hedge fund determines its demand for the risky asset as a function of price. The net demand of the noise traders is then realized and trading occurs, which determines the equilibrium price and quantity traded for the risky asset. In period 2, the dividend on the risky asset is realized and paid.

A.1: The Hedge Fund's Objective Function

The objective function of the hedge fund is assumed to take a simple quadratic form, and the fund is subject to one additional, important constraint. In our model, the hedge fund is concerned

about the illiquidity of its portfolio, because it potentially faces outflows from its investors. As a result, it requires a sufficient quantity of liquid assets in each period to pay out investors who withdraw their funds.

We model the illiquidity constraint in a straightforward fashion. The hedge fund has an expectation of the maximum level of fund outflows in a given time period, denoted $\Phi_{\max}$, which we assume to be increasing in both C_0 and θ. We model a convex "shortfall penalty" if the fund does not hold sufficient liquid assets to cover $\Phi_{\max}$ once it trades in the risky asset. Letting x_1 denote the demand of the hedge fund for the risky asset, and p_1 the equilibrium price of the risky asset in period 1, the convex shortfall penalty is given by

$$
\begin{array}{ll}
0 & \text{if} \quad \Phi_{\max} < C_1, \\[2mm]
\frac{1}{2}(\Phi_{\max} - C_1)^2 & \text{if} \quad \Phi_{\max} \geq C_1,
\end{array}
\tag{A.1}
$$

where $C_1 = C_0 - x_1 p_1$ is the new level of liquid assets that the fund holds after its purchases or sales of the risky asset.

The fund's objective function can therefore be written as

$$
Q(x_1) = x_1(E[d_2] - p_1) - \frac{\psi}{2}x_1^2\sigma_d^2 - \lambda\frac{1}{2}(\Phi_{\max} - C_0 + x_1 p_1)^2 \mathbf{1}\{\Phi_{\max} \geq C_0 - x_1 p_1\}. \tag{A.2}
$$

Here, d_2 is the dividend, which has variance σ_d^2, ψ is the risk aversion of the fund, and $\lambda > 0$ is the weight on the illiquidity constraint in the fund's objective function.

The first-order condition of the hedge fund's maximization problem is[16]

$$
(E[d_2] - p_1) - \psi x_1\sigma_d^2 - \lambda(\Phi_{\max} - C_0 + x_1 p_1)p_1\mathbf{1}\{\Phi_{\max} \geq C_0 - x_1 p_1\} = 0, \tag{A.3}
$$

This results in the period 1 demand function

$$
x_1 = \frac{(E[d_2] - p_1) - \lambda(\Phi_{\max} - C_0)p_1\mathbf{1}\{\Phi_{\max} \geq C_0 - x_1 p_1\}}{\psi\sigma_d^2 + \lambda p_1^2\mathbf{1}\{\Phi_{\max} \geq C_0 - x_1 p_1\}}. \tag{A.4}
$$

The demand function reveals that the lower the value of C_0, and the higher the illiquid endowment θ (which raises $\Phi_{\max}$), the lower is the quantity of the risky asset x_1 that the hedge fund will hold. This yields an important prediction: an illiquid hedge fund is more willing to sell the risky asset than to buy it. This feature of the model, which we explain more fully below, is what underpins the model's explanation for the predictive power of our hedge fund illiquidity index for asset returns.

A.2: Market Clearing and the Equilibrium Price

The demand of the noise traders u_1 is an exogenous shock, and determines the price that the hedge fund demands to absorb this shock in equilibrium. To compute the equilibrium price we need to clear markets. As in Gromb and Vayanos (2010), we interpret the demand shock u_1 as net aggregate demand, which implies that the risky asset is in zero net supply. The market clearing condition therefore takes the form

$$
\frac{(E[d_2] - p_1) - \lambda(\Phi_{\max} - C_0)p_1\mathbf{1}\{\Phi_{\max} \geq C_0 - x_1 p_1\}}{\psi\sigma_d^2 + \lambda p_1^2\mathbf{1}\{\Phi_{\max} \geq C_0 - x_1 p_1\}} + u_1 = 0. \tag{A.5}
$$

---

[16] We note that the first derivative exists for every x_1 in the domain. A proof that the first derivative exists at the jump of the indicator function x_1^o, where $\Phi_{\max} - C_0 + x_1^o p_1 = 0$, can be found in the Internet Appendix.

For the case in which $\Phi_{\max} \geq C_0 - x_1 p_1$, solving for p_1 results in the following equilibrium price[17]

$$p_1 = \frac{(1 + \lambda(\Phi_{\max} - C_0)) - \sqrt{(1 + \lambda(\Phi_{\max} - C_0))^2 - 4(\lambda u_1)(E[d_2] + \psi \sigma_d^2 u_1)}}{2\lambda u_1}. \quad (A.6)$$

For $\Phi_{\max} < C_0 - x_1 p_1$, the equilibrium price is

$$p_1 = E[d_2] + u_1 \psi \sigma_d^2. \quad (A.7)$$

A.3: Model-Implied Hedge Fund Illiquidity and Asset Returns

We measure the model-implied illiquidity level of the hedge fund's portfolio at time 1 as the ratio of the value of its holdings of illiquid assets to the value of its total assets under management before trading the risky asset:

$$\rho_1^{Model} = \frac{\theta}{\theta + C_0}. \quad (A.8)$$

Hence, the lower C_0 relative to θ, the greater the illiquidity of the hedge fund's portfolio. The implicit assumption here is that the illiquid endowment and the risky asset contribute equally to illiquidity–an assumption that can of course be relaxed; we look at comparative statics along this dimension when testing model-implied predictions.

Next, we specify parameters and simulate the model to better understand its predictions. For this analysis, we set $\Phi_{\max} = \phi_{\max}(\theta + C_0)$. For example, when $\phi_{\max} = 0.5$, the hedge fund expects that in the worst-case scenario its investors will withdraw 50% of the AUM in the next period.

$\phi_{\max}$ is determined endogenously. We set $\phi_{\max} = (\theta/(\theta + C_0))^2$. The intuition is that the higher the ratio of the illiquid endowment to the total AUM, the more concerned investors become about the illiquidity of the hedge fund. Due to stronger investor illiquidity concerns, redemptions are expected to be higher. We assume a convex relationship between portfolio illiquidity and redemptions, although a linear relationship would lead to similar results.

Panel A of Table A.1 shows the specific parameter values that we use. The starting endowment of cash is set to 4. The two sources of uncertainty in our model, d_2 and u_1, are drawn from Normal distributions with means $E[d_2] = 1$ and $E[u_1] = 0$, and variances $\sigma_d^2 = 0.002$ and $\sigma_u^2 = 0.25$, respectively. Different parameter values would lead to predictions of the model that are qualitatively the same. However, $\sigma_d^2 = 0.002$ is chosen to match the empirical variance of monthly U.S. equity excess returns. The variance of u_1 is more difficult to pin down empirically. However, a value of 0.25 causes the hedge fund in our model to use on average around 3% of its AUM for liquidity provision, which seems to be a reasonable assumption. λ, the weight on the shortfall penalty in the hedge fund's objective function, is 0.002, and ψ, the risk-aversion of the hedge fund, is set to 3.[18] Due to the two regimes shown in Equation (A.1), we need a numerical procedure to simulate the model. The numerical procedure is described in detail in Section A.4.

Panel B of Table A.1 looks at comparative statics as we vary θ, which determines the level of illiquidity of the fund. In this exercise, values for d_2 and u_1 are drawn from their distributions whose parameters are given in panel A, while the value of θ is varied across different simulations. We simulate the two-period model 10,000 times for each value of θ and compute the average expected return, ρ_1^{Model}, and the "price discounts" (relative to the price when $\lambda = 0$) at which noise

---

[17] The quadratic equation for the equilibrium price has a second root. However, we ignore the second root, as the implied equilibrium price is not sensible for the parameter values in our model simulation.

[18] The choice of λ is discussed in more detail below.

Table Appendix A.1
Model simulation

Panel A: Parameter values

C_0	λ	ψ	$E[d_2]$	σ_d^2	$E[u_1]$	σ_u^2
4	0.002	3	1	0.002	0	0.25

Panel B: Model simulation for fixed θ

No liquidity constraint

θ	C_0/θ	$\phi_{max}(\theta + C_0)$	ρ_1^{Model}	Expected return	Noise traders buy price	Noise traders sell price
na	na	na	na	0.000%	1.002	0.998

With liquidity constraint

θ	C_0/θ	$\phi_{max}(\theta + C_0)$	ρ_1^{Model}	Expected return	Price discount NT buy	NT sell
8.0	0.5	5.3	0.667	0.266%	0.187%	−0.345%
10.0	0.4	7.1	0.714	0.624%	0.546%	−0.702%
12.0	0.3	9.0	0.750	0.992%	0.911%	−1.066%

Panel C: Predictive regression on simulated data

	Parameter estimates	
	(1)	(2)
α^{Model}	−0.016	−0.012
γ^{Model}	0.033	
$\gamma^{Model, NT\ Buy}$		0.028
$\gamma^{Model, NT\ Sell}$		0.037

Panel A reports the parameter values used for the model simulations in panel B and C. Panel B reports the statistics for model simulations based on fixed values of θ. For each value of θ, the model is simulated 10,000 times. The parameters d_2 and u_1 are drawn randomly from Normal distributions with means of 1 and 0 (variances 0.002 and 0.25, respectively, for each of the 10,000 simulations. The ρ_1^{Model}, returns, and prices are averaged across the 10,000 simulations. Panel C reports the estimates of α^{Model} and γ^{Model} for the predictive regression $r_2 = \alpha^{Model} + \gamma^{Model} \rho_1^{Model} + \epsilon_2$ estimated with 10,000 data points of simulated data from our model. The values for d_2, u_1, and θ are drawn randomly for each of the 10,000 simulations. The values for d_2 and u_1 are drawn from the same distributions as in panel B. θ is drawn from a uniform distribution with a lower bound of 0 and an upper bound of 20. The value for C_0 is 4.

traders buy from and sell to the hedge fund. Since the form of the demand function depends on $\Phi_{max} \geq C_0 - x_1 p_1$, we need a numerical procedure to simulate the model. The procedure is shown in Appendix A.4.

The top row in this panel simply sets the liquidity constraint on the fund to zero, that is, $\lambda = 0$, as a benchmark, so the values of θ and ϕ_{max} are not relevant. The expected return is computed as the average realization of $\frac{1}{p_1}$. The noise traders buy from the hedge fund at a price slightly above expected fundamental value and sell to the hedge fund at a slightly lower price than expected fundamental value as compensation for the hedge fund's market-making services.

As we move down the rows of panel B of the table, we impose the liquidity constraint, and vary the level of the illiquid endowment, θ, resulting in a cash-to-liquid-assets ratio of 50%, 40%, and 33%, respectively, for each of the bottom three rows. The ρ_1^{Model} values that these correspond to are shown in the fourth column, and we can see that an increase in θ leads to a higher average
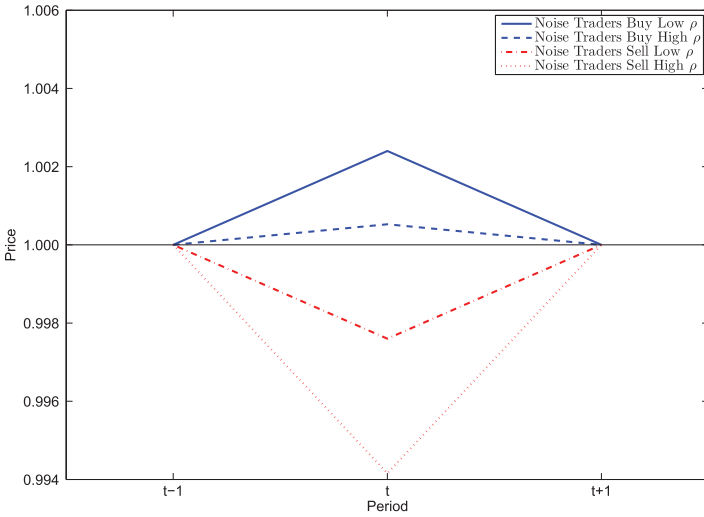
Figure Appendix A.1
Model return reversal: This figure depicts the average return reversal implied by the model for the parameter values given in Table Appendix A.1. The high ρ case corresponds to $\theta = 8$ and $\lambda = 0.002$, and the low ρ case corresponds to the model without shortfall penalty, that is, $\lambda = 0$.

ρ_1^{Model}, as well as in the next column, to a higher average expected return. This expected return is higher than the return in the benchmark case, which reflects the extra compensation needed to convince the illiquid hedge fund to purchase more of the risky asset, as well as the price discount offered by the hedge fund to the noise traders when they purchase the risky asset. While the illiquid hedge fund is willing to sell the risky asset for a lower price in the face of a positive noise trader demand shock, to increase the liquidity of its portfolio and avoid hitting the shortfall penalty, the fund is reluctant to buy the risky asset and requires more compensation to do so. Relative to expected fundamental value, the price discount to the noise traders is substantial given our parameter values–reaching roughly 0.9% when the hedge fund discounts the price to sell to noise traders, and roughly -1.1% when the hedge fund demands compensation in order to accept more of the risky asset into its portfolio, when $\theta = 12$.

Finally, we again simulate the model 10,000 times, with θ drawn from a uniform distribution with a lower bound of 0 and an upper bound of 20, and d_2 and u_1 drawn as described above. We then take the 10,000 observations of ρ_1^{Model} derived from each one of these simulations and the realized returns in the second period r_2, derived from these simulations, and regress:

$$r_2 = \alpha^{Model} + \gamma^{Model} \rho_1^{Model} + \epsilon_2. \tag{A.9}$$

The model-implied α^{Model} and γ^{Model} are reported in the first column of panel C of Table A.1. We also estimate separate coefficients on ρ_1^{Model} for noise traders buying and selling. The results are shown in Column 2. They reveal that the coefficient γ is positive regardless of the sign of the shock, which is consistent with the intuition of the model, and a feature that we confirm in our auxiliary empirical tests. We can see that $\gamma^{Model\ NT\ Buy}$ is smaller than $\gamma^{Model\ NT\ Sell}$. The reason is that the shortfall penalty is less likely to be triggered when noise traders buy than when noise traders sell. This can be seen in Equation (A.1).

Figure A.1 provides a graphical representation of the model predictions: the figure shows the comparative static of the price path compared across levels of illiquidity ρ_1^{Model} of the hedge fund. Compared to a liquid hedge fund, an illiquid hedge fund buys for a *lower* price when noise traders sell, implying a *greater* return reversal. It also sells for a *lower* price when noise traders buy the

Table Appendix A.2
λ sensitivity of model R^2

λ	Model R^2	Avg. price	SD ρ_1^{Model}	α^{Model}	γ^{Model}	Shortfall elasticity of demand
0.0005	0.779%	0.994	0.086	−0.031	0.047	−0.781%
0.0010	2.720%	0.987	0.085	−0.060	0.091	−1.435%
0.0015	5.749%	0.981	0.087	−0.084	0.129	−2.183%
0.0020	9.878%	0.975	0.087	−0.118	0.179	−2.962%

λ, the model R^2, the average price, the standard deviation of ρ_1^{Model}, and the estimates of α^{Model} and γ^{Model} for the predictive regression $r_2 = \alpha^{Model} + \gamma^{Model} \rho_1^{Model} + \epsilon_2$ estimated with 10,000 data points of simulated data from our model are reported. The last column presents the average shortfall elasticity of demand across the price range from 0.95 to 1.05 (for the expected value of θ): the percentage change in quantity of the risky asset demanded based on a 1% change in the cash short fall, $\Phi_{max} - C_0$. The values for d_2, u_1, and θ are drawn randomly for each of the 10,000 simulations. The values for d_2 and u_1 are drawn from Normal distributions with means of 1 and 0 (variances 0.002 and 0.25), respectively, for each of the 10,000 simulations. θ is drawn from a uniform distribution with a lower bound of 5 and an upper bound of 35. C_0 is set equal to 4, and the value of ψ is 3.

risky asset, implying a *smaller* return reversal in absolute values. When these predictions are taken together, the model implies that high ρ_1^{Model} predicts high returns.

The model also provides some insight into the R^2 we can expect in our empirical regressions. The parameter λ, which measures the sensitivity of the hedge fund to a shortfall in cash, is difficult to match empirically but has a big impact on the predictive power of ρ_1^{Model}. Table A.2 shows the sensitivity of the Model R^2 from regression (A.9) to λ. The parameter values are the same as used in Table A.1, panel C, except that θ is drawn from a uniform distribution with a lower bound of 5 and an upper bound of 35 to match the empirical standard deviation of ρ, which is 0.085. An increase in λ leads to the hedge fund reacting more sensitively to a cash shortfall, which results in a lower average price for the risky asset and stronger predictive power of ρ_1^{Model}. Correspondingly, the γ^{Model} estimate also increases. The empirical value of λ is determined by a range of factors, as, for example, margin requirements and illiquidity of the risky asset. This explains why we find that ρ has more predictive power during the financial crisis as shown in Figure 3: funding constraints make hedge funds more sensitive to a cash shortfall, and thus, the predictive power of ρ is greater. Our simple model does not account for margin requirements or variations in the illiquidity of the risky asset. However, the empirical predictability of ρ that we find can be closely matched by the model R^2 for a range of λ values. For example, the predictive power of ρ for Australia's equity index is 7.487% and matched by the model with a λ between 0.0015 and 0.0020, but the empirical $\overline{R}^2$ for the United States is 1.027% and closely matched by the model with a λ of 0.0005. The last column of Table A.2 shows the average shortfall elasticity of demand for these λ values. The elasticity is averaged across the price range from 0.95 to 1.05 (for the expected value of θ) and is defined as the percentage change in quantity of the risky asset demanded based on a 1% change in the cash shortfall, $\Phi_{max} - C_0$. We see that the greater λ, the more elastic the demand becomes.

A.4: Numerical Procedure for Model Simulation

As the form of the demand function depends on x_1 due to the indicator function, we need a numerical procedure to simulate the model. The procedure has three steps:

1. Use the demand function $x_1 = \frac{E(d_2) - p_1}{\psi \sigma_d^2}$ and determine p_1^* at which the noise trader demand $u_1^* = -x_1^*$ clears.
2. Check whether $\Phi_{max} < C_0 - x_1^* p_1^*$. If the inequality holds, the solution is found. If the inequality does not hold, proceed with 3.
3. Use the demand function $x_1 = \frac{E(d_2) - p_1 - \lambda(\Phi_{max} - C_0)p_1}{\psi \sigma_d^2 + \lambda p_1^2}$, and determine $\hat{p}_1$ at which the noise trader demand $\hat{u}_1 = -\hat{x}_1$ clears. $\Phi_{max} \geq C_0 - \hat{x}_1 \hat{p}_1$ will hold conditional on $\Phi_{max} \geq C_0 - x_1^* p_1^*$.

Appendix B: A Modification of Clark and West (2006)

Clark and West (2006) consider the problem of testing equal predictive accuracy of a linear regression and a model with no parameters. In this Appendix we propose a simple modification that allows the "smaller" model to include a constant. (Clark and West 2007 extend their 2006 paper to consider cases in which the smaller model includes predictor variables, but they do not obtain asymptotic critical values for their testing approach. Our focus on the case that the smaller model includes *only* a constant means our problem is simpler, and leads us to a simple testing procedure.)

Denote the two forecasts as

$$\hat{Y}^{(1)}_{t+1|t} = \hat{\gamma}_t,$$
$$\hat{Y}^{(2)}_{t+1|t} = X'_{t+1}\hat{\beta}_t. \tag{B.1}$$

The second forecast is based on a set of predictor variables X_{t+1}, which includes a constant, and so nests the smaller model. Both forecasts are based on parameters estimated using a rolling window of data (and so they are not constant through the sample) of fixed length R. The sample period runs from $t = 1, 2, \ldots, R, R+1, \ldots, R+P+1 \equiv T$. The mean squared prediction errors (MSPEs) over the out-of-sample period ($t = R+1, \ldots, R+P$) are

$$\hat{\sigma}^2_1 = \frac{1}{P}\sum_{t=R+1}^{R+P}(y_{t+1} - \hat{\gamma}_t)^2,$$
$$\hat{\sigma}^2_2 = \frac{1}{P}\sum_{t=R+1}^{R+P}(y_{t+1} - X'_{t+1}\hat{\beta}_t)^2. \tag{B.2}$$

The null and alternative hypotheses are

$$H_0 : \sigma^2_1 = \sigma^2_2 \text{ vs. } H_1 : \sigma^2_1 > \sigma^2_2, \tag{B.3}$$

That is, the small and large models are equally good under the null, while under the alternative the small model is worse than the large model. Note that under the null, where the smaller model is correct, we have

$$y_{t+1} = \gamma^* + \varepsilon_{t+1}, \, E_t[\varepsilon_{t+1}] = 0. \tag{B.4}$$

The difference in the MSPEs of these two forecasts is

$$\hat{\sigma}^2_1 - \hat{\sigma}^2_2 = \frac{1}{P}\sum_{t=R+1}^{R+P}\left\{(y_{t+1} - \hat{\gamma}_t)^2 - (y_{t+1} - X'_{t+1}\hat{\beta}_t)^2\right\}$$
$$= \frac{1}{P}\sum_{t=R+1}^{R+P}\hat{\gamma}^2_t - 2\frac{1}{P}\sum_{t=R+1}^{R+P}y_{t+1}\hat{\gamma}_t - \frac{1}{P}\sum_{t=R+1}^{R+P}(X'_{t+1}\hat{\beta}_t)^2 + 2\frac{1}{P}\sum_{t=R+1}^{R+P}y_{t+1}X'_{t+1}\hat{\beta}_t \tag{B.5}$$

and so

$$\hat{\sigma}^2_1 - \hat{\sigma}^2_2 \xrightarrow{p} E[\hat{\gamma}^2_t] - E[(X'_{t+1}\hat{\beta}_t)^2] - 2\gamma^* E[(\hat{\gamma}_t - X'_{t+1}\hat{\beta}_t)] \text{ under } H_0 \text{ as } P\to\infty \tag{B.6}$$

using the fact that $E[\varepsilon_{t+1}X'_{t+1}\hat{\beta}_t] = E[\varepsilon_{t+1}\hat{\gamma}_t] = 0$.

When $\hat{\gamma}_t = 0 \forall t$ and $\gamma^* = 0$, we are in the Clark-West framework, and they note that

$$\hat{\sigma}^2_1 - \hat{\sigma}^2_2 \xrightarrow{p} -E[(X'_{t+1}\hat{\beta}_t)^2] \text{ under } H_0 \text{ as } P\to\infty, \tag{B.7}$$

That is, when the smaller model is correct, the difference in MSPEs will be centered on a *negative value*, even though the two models are both correct (the larger model nests the smaller model, so it is also correct). This makes the test conservative (since we use a standard Normal as the

asymptotic distribution of the test statistic, which is centered on zero) and so it will be hard to reject in favor of the larger model when it is correct. Clark and West suggest adjusting $\hat{\sigma}_2^2$ so that the difference between the MSPEs is centered on zero under the null:

$$\hat{\sigma}_{2,adj}^2 = \hat{\sigma}_2^2 - \frac{1}{P} \sum_{t=R+1}^{R+P} (X_{t+1}' \hat{\beta}_t)^2. \tag{B.8}$$

They show that this provides better properties under the null and under the alternative.

We extend this adjustment to allow the smaller model to include a constant. Intuitively, allowing for a constant adds a small amount of variability to forecast 1, making the extra penalty faced by forecast 2 slightly lower in relative terms. It also introduces some cross-product terms that need to be handled. Let us define an adjusted difference in MSPEs to correct for these terms:

$$\hat{\Delta}_{adj} = \hat{\sigma}_1^2 - \hat{\sigma}_2^2 - \frac{1}{P} \sum_{t=R+1}^{R+P} \hat{\gamma}_t^2 + \frac{1}{P} \sum_{t=R+1}^{R+P} (X_{t+1}' \hat{\beta}_t)^2$$
$$+ 2\hat{\gamma}^* \frac{1}{P} \sum_{t=R+1}^{R+P} (\hat{\gamma}_t - X_{t+1}' \hat{\beta}_t), \tag{B.9}$$

where

$$\hat{\gamma}^* = \frac{1}{P} \sum_{t=R+1}^{R+P} y_{t+1}. \tag{B.10}$$

Next, we need to get a limiting distribution for $\hat{\Delta}_{adj}$. In the Clark-West case $\hat{\Delta}_{adj}$ is just a linear combination of sample averages, and so this can be obtained by defining an "adjusted difference in per-period loss variable" ($\hat{f}_{t+1}$ in their equation 3.3), and conducting a t-test that that variable is zero mean. In our case, the cross-product term introduces a *product* of sample averages, and so we cannot use that approach. Instead, we use the delta method to get the limiting distribution of our test statistic. Let

$$\underset{(6 \times 1)}{g_t} = [(y_{t+1} - \hat{\gamma}_t)^2, -(y_{t+1} - X_{t+1}' \hat{\beta}_t)^2, -\hat{\gamma}_t^2, (X_{t+1}' \hat{\beta}_t)^2, 2(\hat{\gamma}_t - X_{t+1}' \hat{\beta}_t), y_{t+1}]'. \tag{B.11}$$

Under standard regularity conditions, we obtain

$$\sqrt{P}(\bar{g}_P - \bar{g}_0) \xrightarrow{d} N(0, V_g) \text{ as } P \to \infty. \tag{B.12}$$

As in Clark and West (2006), the asymptotic covariance matrix, V_g, should be estimated using a HAC estimator (for example, Newey-West). Our test statistic is a nonlinear function of $\bar{g}_P$

$$\hat{\Delta}_{adj} = f(\bar{g}_P) = \bar{g}_1 + \bar{g}_2 + \bar{g}_3 + \bar{g}_4 + \bar{g}_5 \bar{g}_6 \tag{B.13}$$

and so by the delta method we obtain

$$\sqrt{P}\big(f(\bar{g}_P) - f(\mathbf{g}_0)\big) = \sqrt{P}(\hat{\Delta}_{adj} - \Delta_0) \xrightarrow{d} N\big(0, \nabla_g f(\bar{g}_P) V_g \nabla_g f(\bar{g}_P)'\big), \tag{B.14}$$

where $\nabla_g f(\bar{g}_P) = \partial f(\bar{g}_P)/\partial \mathbf{g}' = [1, 1, 1, 1, \bar{g}_6, \bar{g}_5]$.

We can use this to obtain a Clark-West style test for the larger model versus a model just including a constant. Specifically, we compute the statistic $\sqrt{P}\hat{\Delta}_{adj}/\sqrt{\partial f(\bar{g}_P)\hat{V}_g \partial f(\bar{g}_P)'}$ and compare it to the right-tail critical values of the $N(0, 1)$ distribution (1.28, 1.65, 2.33) to get a test at the 10%, 5%, or 1% level.

The Internet Appendix presents a small simulation study verifying that this test has satisfactory finite-sample properties and confirming that it also leads to power gains relative to the unadjusted test.

References

Acharya, V., and L. H. Pedersen. 2005. Asset pricing with liquidity risk. *Journal of Financial Economics* 77:375–410.

Agarwal, V., N. D. Daniel, and N. Y. Naik. 2011. Do hedge funds manage their reported returns? *Review of Financial Studies* 24:3281–320.

Agarwal, V., and N. Y. Naik. 2004. Risks and portfolio decisions involving hedge funds. *Review of Financial Studies* 17:63–98.

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

Aragon, G. O., B. Liang, and H. Park. 2014. Onshore and offshore hedge funds: Are they twins? *Management Science* 60:74–91.

Aragon, G. O., and P. E. Strahan. 2012. Hedge funds as liquidity providers: Evidence from the Lehman bankruptcy. *Journal of Financial Economics* 103:570–87.

Ben-David, I., F. Franzoni, and R. Moussawi. 2012. Hedge fund stock trading in the financial crisis of 2007–2009. *Review of Financial Studies* 25:1–54.

Bollen, N. P. B., and V. K. Pool. 2008. Conditional return smoothing in the hedge fund industry. *Journal of Financial and Quantitative Analysis* 43:267–98.

———. 2009. Do hedge fund managers misreport returns? Evidence from the pooled distribution. *Journal of Finance* 64:2257–88.

Bollen, N. P. B., and R. E. Whaley. 2009. Hedge fund risk dynamics: Implications for performance appraisal. *Journal of Finance* 64:987–1037.

Bongaerts, D., F. de Jong, and J. Driessen. 2012. An asset pricing approach to liquidity effects in corporate bond markets. Unpublished Working Paper.

Brunnermeier, M., and L. H. Pedersen. 2009. Market liquidity and funding liquidity. *Review of Financial Studies* 22:2201–38.

Burnside, C., M. Eichenbaum, and S. Rebelo. 2011. Carry trade and momentum in currency markets. *Annual Review of Financial Economics* 3:511–35.

Campbell, J. Y.. 1991. A Variance decomposition for stock returns. *The Economic Journal* 101:157–79.

Campbell, J. Y., S. J. Grossman, and J. Wang. 1993. Trading volume and serial correlation in stock returns. *Quarterly Journal of Economics* 108:905–39.

Campbell, J. Y., K. Serfaty-de Medeiros, and L. M. Viceira. 2010. Global currency hedging. *Journal of Finance* 65:87–122.

Campbell, J. Y., and S. B. Thompson. 2008. Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies* 21:1509–31.

Cao, C., Y. Chen, W. N. Goetzmann, and B. Liang. 2015. The role of hedge funds in the security price formation process. Unpublished Working Paper.

Cao, C., Y. Chen, B. Liang, and A. Lo. 2013. Can hedge funds time market liquidity? *Journal of Financial Economics* 109:493–516.

Clark, T. E., and K. D. West. 2006. Using out-of-sample mean squared prediction errors to test the martingale difference hypothesis. *Journal of Econometrics* 135:155–86.

———. 2007. Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics* 138:291–311.

Dimson, E. 1979. Risk measurement when shares are subject to infrequent trading. *Journal of Financial Economics* 7:197–226.

Ferson, W. E., S. Sarkissian, and T. Simin. 2003. Is stock return predictability spurious? *Journal of Investment Management* 1:1–10.

Fung, W., D. A. Hsieh. 1997. Empirical characteristics of dynamic trading strategies: The case of hedge funds. *Review of Financial Studies* 10:275–302.

———. 2001. The Risk in hedge fund strategies: Theory and evidence from trend followers. *Review of Financial Studies* 14:313–41.

———. 2004. Hedge fund benchmarks: A risk based approach. *Financial Analyst Journal* 60:65–80.

Fung, W., D. A. Hsieh, N. Y. Naik, and T. Ramadorai. 2008. Hedge funds: Performance, risk, and capital formation. *Journal of Finance* 63:1777–803.

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

Goyal, A., and I. Welch. 2008. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies* 21:1455–1508.

Gromb, D., and D. Vayanos. 2002. Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics* 66:361–407.

———. 2010. Limits of arbitrage: The state of the theory. *Annual Review of Financial Economics* 2:251–75.

Jylha, P., K. Rinne, and M. Suominen. 2014. Do hedge funds supply or demand liquidity? *Review of Finance* 18:1259–98.

Kang, N., P. Kondor, and R. Sadka. 2014. Do hedge funds reduce idiosyncratic risk? *Journal of Financial and Quantitative Analysis* 49:843–77.

Khandani, A. E., and A. W. Lo. 2011. What happened to the quants in August 2007? Evidence from factors and transactions data. *Journal of Financial Markets* 14:1–46.

Lo, A. W. 2008. Hedge funds: An analytic perspective. Princeton: Princeton University Press.

Meese, R. A., and K. Rogoff. 1983. Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics* 24:3–24.

Mitchell, M., L. H. Pedersen, and T. Pulvino (2007). Slow moving capital. *American Economic Review P&P* 97:215–20.

Newey, W. K., and K. D. West. 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55:703–08.

Pastor, L., and R. Stambaugh. 2003. Liquidity risk and expected stock returns. *Journal of Political Economy* 111:642–85.

Patton, A. J., and T. Ramadorai. 2013. On the high-frequency of hedge fund risk exposures. *Journal of Finance* 68:597–635.

Patton, A. J., T. Ramadorai, and M. Streatfield. 2015. Change you can believe in? Hedge fund data revisions. *Journal of Finance* 70:963–99.

Sadka, R. 2010. Liquidity risk and the cross-section of hedge-fund returns. *Journal of Financial Economics* 98:54–71.

———. 2011. Liquidity risk and accounting information. *Journal of Accounting and Economics* 52:144–52.

Scholes, M., and J. Williams. 1977. Estimating betas from nonsynchronous data. *Journal of Financial Economics* 5:309–27.

Shao, J. 1989. The efficiency and consistency of approximations to the jackknife variance estimators. *Journal of the American Statistical Association* 84:114–9.

Shao, J., and C. F. J. Wu. 1989. A general theory for jackknife variance estimation. *Annals of Statistics* 17:1176–97

Shleifer, A., and R. W. Vishny. 1997. The limits of arbitrage. *Journal of Finance* 52:35–55.

Stambaugh, R. J. 1999. Predictive regressions. *Journal of Financial Economics* 54:375–421.