Hedge Fund Predictability Under the Magnifying Glass: The Economic Value of Forecasting Individual Fund Returns*

Doron Avramov†, Laurent Barras‡, and Robert Kosowski§

First version, June 5th, 2008;
This version, October 31st, 2010

JEL Classification: G11, G23, C12
Keywords: Hedge Fund Performance, Return Predictability, Combination Forecasts

ABSTRACT

The recent financial crisis has highlighted the need to search for suitable models predicting hedge fund performance. This paper embraces the challenge by developing and applying a unified framework in which to comprehensively assess predictability across individual funds. In-sample, we identify the fraction of hedge funds in the overall universe as well as distinct investment styles that are truly predictable, positively or negatively, by macro variables. Out-of-sample, we show that exploiting predictability in a multi-fund environment is challenging as strongly predictable funds typically exhibit low unconditional means. Moreover, estimation risk and model instability lead to imprecise return forecasts. Nevertheless, combining forecasts across predictors circumvents all these difficulties and delivers superior performance. We highlight the economic and statistical drivers of this performance, and examine the 2008 crisis, when predictor values strongly depart from their long run means, as a natural experiment to validate the robustness of our methodology.

*We thank Bill Fung, Tarun Ramadorai, Olivier Scaillet, Stas Sokolinski, Melvyn Teo, Jialin Yu, as well as seminar participants at HEC-Montreal, Imperial College, McGill University, Tilburg University, the University of Rotterdam, the University of Toronto, the 4th Imperial College London conference on hedge funds, the 2nd INSEE/CREST conference on hedge funds, the 2009 Annual Meeting of the European Finance Association (EFA), and the 2010 Annual Meeting of the Institute for Mathematical Finance (IFM2) for their comments.
†Hebrew University of Jerusalem and University of Maryland. email: davramov@rhsmith.umd.edu
‡Desautels Faculty of Management, McGill University. email: laurent.barras@mcgill.ca
§Imperial College Business School, Imperial College London. email: r.kosowski@imperial.ac.uk

I Introduction

During the recent financial crisis, the hedge fund industry suffered its worst performance ever. While predicting crises is inherently difficult, traditional measures building on past stellar performance conclusively failed. For one, the formerly best performing styles, long-short equity and emerging markets, recorded a 13.8% and 29.3% loss in 2008. Chicago-based Citadel, the $11 billion investment firm founded by Ken Griffin, exhibits one fund specific example that investment opportunities do depend on the state of the economy. Citadel posted a mid-year 21% return in 2009 after betting that outflows from convertible bonds accompanied by price drops were excessive at the beginning of 2009.[1] Such events raise essential questions that we attempt to address here. Can future hedge fund performance be predicted by variables that capture changing business conditions? If yes, (i) does predictability vary across as well as within distinct investment styles, and (ii) why do certain economic variables better forecast returns for distinct investment styles? Moreover, can hedge fund investors successfully exploit predictability out-of-sample to improve performance? If yes, (i) what are the economic and statistical drivers of this superior performance, and (ii) what is the forecast quality during periods when predictor values dramatically depart from their long run levels, as in 2008?

This paper attempts to resolve these questions. In particular, it develops and applies a unified econometric framework in which to comprehensively analyze predictability, both in- and out-of-sample, at the individual fund level. The fund level analysis is essential because hedge funds follow a wide range of investment strategies, even within pre-established investment styles. While under specific economic conditions some particular funds perform well, others do poorly. Hence, studying individual funds rather than broad indices, is more informative, as cross-fund differences may simply average out. Moreover, any realistic portfolio advice that exploits predictability amounts to selecting individual funds, as broad hedge fund indices are not investable.

To precisely assess the extent of in-sample predictability, we determine the proportion of funds in the population having future returns that are (i) negatively related; (ii) unrelated; or (iii) positively related to any given predictive variable. Importantly, this approach can distinguish between true and spurious predictability, and allows for a careful analysis of predictability both within and across investment styles.

To exploit predictability out-of-sample, we implement a conditional strategy that selects, at each rebalancing date, the top decile of funds exhibiting the highest conditional mean. Specifically, we use a *combination* strategy in which the fund return forecasts,

[1] See "Canyon, Citadel Ride Convertibles to Recoup Losses", Bloomberg article, June 11, 2009.

obtained from each predictive variable separately, are averaged to estimate the conditional mean. This strategy features several advantages. First, because the conditional mean is the sum of the unconditional and predictable performance, the combination strategy uses predictive information parsimoniously–only when the predictor values depart from their long run level. This is important because in a large fund population, funds exhibiting the strongest predictability typically do not deliver the highest unconditional performance. In addition, the combination strategy accounts for estimation risk and model uncertainty. Indeed, the identity of predictable funds must be estimated based on data, thus confronting estimation risk. Model instability arises from policy surprises, institutional changes, or advances in information technology. If the quality of predictive information is poor, even a parsimonious use of predictability may not be economically valuable. By diversifying across forecasts, especially when forecast errors are not fully correlated across models, the combination strategy reduces the impact of poor information quality, just like a portfolio diversifies across assets to reduce risk.

The framework developed here extends past work to address the small sample bias in predictive regressions, accommodating hedge fund benchmarks as well as the return autocorrelation arising from illiquid positions undertaken by hedge funds. Ultimately, our unified approach allows one to establish important links between in- and out-of-sample predictability.

Our empirical analysis covers more than 7,000 hedge funds between 1994 and 2008. The evidence reveals ample new evidence about hedge fund return predictability. Of the overall universe of hedge funds, about 60.5% are predictable using a set of four economically motivated predictors: the default spread, the dividend yield, the VIX, and the net aggregate flows into the hedge fund industry. While each variable taken individually has some predictive ability, we document a strong asymmetry in the direction of this predictability. To illustrate, while 13.4% of the funds are positively predicted by the default spread, only 2.8% have a negative exposure to it. Thus, for the vast majority of predictable funds (83%), times when the default spread is high (low) are associated with higher (lower)-than-average future returns.

Comparing investment categories, we find that the proportion of predictable funds using all predictors varies distinctively, ranging from 31.1% (managed futures) to 83.7% (convertible arbitrage). Consistent with economic rationale, some investment styles exhibit a strong relation with specific predictors. For instance, a high proportion of emerging market funds (33.5%) display future returns that are positively related to the default spread. In fact, widening credit spreads typically coincide with a flight to quality, widening emerging market sovereign bond spreads, and higher future returns.

Another example is the significant proportion of managed futures funds (16.7%) that are positively exposed to the volatility index. Extreme volatility may trigger trend reversals, which is generally beneficial to managed funds. We also observe some common patterns across hedge fund styles in response to changing current values of the dividend yield and net aggregate inflows. In all but one category (managed futures), a high dividend yield signals lower future returns, consistent with more limited access to leverage during recessions. Likewise, across all styles, individual fund returns tend to be negatively related to inflows, consistent with the capacity constraints documented for hedge fund indices by Naik, Ramadorai, and Stromquist (2007). Interestingly, most predictability across all investment styles is attributable to time-varying managerial skills in identifying attractive investment opportunities.

We next analyze the economic value of predictability carrying out a range of out-of-sample tests that incorporate real-world constraints encountered by institutional investors. For one, to account for liquidity constraints (i.e., lock-up periods), we permit only annual portfolio rebalancing. The evidence shows that, over the period 1997-2007, an unconditional portfolio that simply uses past returns to select funds performs well, consistent with Jagannathan, Malakhov, and Novikov (2010). Such unconditional portfolio generates FH alpha, $\widehat{\alpha}$, Information Ratio, IR, and Sharpe Ratio, SR, of 5.8%, 2.4, and 1.8 per year, respectively. For comparison, single-predictor conditional strategies that use only one of the four macro variables to forecast returns underperform the unconditional portfolio on a risk-adjusted basis. Remarkably, the *combination* strategy achieves the highest risk-adjusted performance, producing a FH alpha of 7.0% per year and an additional 50%-cumulative return over the unconditional portfolio. In addition, it exhibits a low tail risk, a low exposure to the FH risk factors (like a "pure alpha" strategy), as well as reasonable levels of turnover and serial correlation.

Why do the single-predictor strategies fail? These strategies exploit predictive information whenever the predictor value considerably moves away, either up or down, from its long run mean. However, in these two cases, the information quality differs due to the asymmetric nature of predictability documented in-sample. To illustrate, consider again the default spread. When its value is below the long run mean, funds with a negative exposure are those that could successfully exploit predictability. Since only 2.8% of the funds in the population have this negative exposure, data will not be of great help to detect such negative exposure funds, leading to poor information quality. Indeed, we find that the out-of-sample performance of the single-predictor strategies perfectly reflects this asymmetry. For instance, in times of a lower-than-average default spread, performance falls substantially.

In contrast, the combination strategy performs extremely well because it is less affected by poor information quality for two reasons. First, it leads to a very conservative portfolio, i.e., 80% of the funds chosen by this strategy are common to the unconditional portfolio. While the benefits of this shrinkage effect are discussed in Rapach, Strauss, and Zhou (2010) in the context of the US stock index, we argue that they are even more important in a multi-asset setting. In such a setting, investors are potentially hit twice–both by choosing funds with low predictability, as well as by excluding funds with high unconditional performance. Second, the diversification across predictors reduces the out-of-sample forecast error variance, and helps detect the truly predictable funds from the data. Indeed, the "active" portfolio chosen by the combination strategy (that is, the remaining 20%) generates a high performance–a FH alpha up to 5.3% per year over the unconditional portfolio–that is not subject to the asymmetry described earlier.

The 2008 financial crisis featured large fluctuations of the predictor values around their long run means, making the predictable signal quite strong. This provides us with a meaningful out-of-sample experiment to measure the economic value of predictive information. Incorporating 2008 reveals that the risk-adjusted performance of the combination strategy is still superior to that of the unconditional portfolio ($\widehat{\alpha} = 6.0\%$ versus 4.1%, $IR = 1.9$ versus 1.2). Among the single-predictor strategies, the VIX strategy resisted remarkably well, posting the lowest final quarter loss among all strategies (3.1% versus 10.8% on average for the other strategies). This positive performance comes, among others, from the higher stability in the predictive relation between the VIX and future hedge fund returns during the crisis. Finally, we examine the cost of the additional liquidity constraints imposed on hedge fund investors to prevent massive outflows during the crisis. Specifically, had investors been able to incorporate predictability at a monthly frequency in 2008, the average annual return across all conditional strategies would have increased from -9.5% up to 3.3%, indicating that investor level illiquidity costs are substantial.

Our paper relates to the vast literature on return predictability. Seminal papers in the field include Keim and Stambaugh (1986), Fama and French (1989), and Ferson and Harvey (1991). On the methodological front, our paper mostly borrows from the literature on combination forecasts (e.g., Bates and Granger (1969), Hendry and Clements (2004), Timmermann (2006)), while extending the Amihud, Hurvich, and Wang (2008) method to correct for small sample bias in predictive regressions. Focusing on hedge fund return predictability, Amenc, El Bied, and Martellini (2003) examine predictability using broad-based hedge fund indices, while Avramov, Kosowski, Naik, and Teo (henceforth AKNT) (2010) form hedge fund portfolios that exploit predictability using a Bayesian framework. While AKNT show that a Bayesian investor could exploit predictability to improve asset allocation decisions, they do not shed light on how and why certain predictors forecast fund returns. Ultimately, we provide new insights about predictability at the fund level, both within and across investment categories. Moreover, the simplicity of the conditional strategies proposed here allows a clear link between in-sample and out-of-sample predictability, as well as establishing a deeper understanding of the economic and statistical drivers of out-of-sample performance.

The paper proceeds as follows. Section II discusses the methodology. Section III describes the data. Section IV summarizes the findings, while Section V concludes.

II Understanding Hedge Fund Predictability

A Predicting Hedge Fund Performance

We attempt to predict future returns on M individual hedge funds using J aggregate variables that potentially capture evolving economic conditions. Predictability is analyzed based on the time series predictive regression, run separately for each fund,

$$r_{i,t+1} = b_{i,0} + \sum_{j=1}^{J} b_{i,j} \cdot Z_{j,t} + u_{i,t+1}. \tag{1}$$

The dependent variable $r_{i,t+1}$ denotes the time $t+1$ excess hedge fund return (over the riskfree rate), $Z_{j,t}$ ($j = 1, ..., J$) is the time t realized value of the j-th predictive variable, $b_{i,0}$ is the intercept, $b_{i,j}$ is the slope coefficients associated with each predictor, and $u_{i,t+1}$ denotes the unpredictable fund specific innovation.

Hedge funds typically follow a wide range of strategies and trade many different assets. As a result, some funds are likely to perform better under specific economic conditions, while others will do poorly. This heterogeneity, captured by cross-fund variation in the predictive regression slope coefficients, provides a strong motivation to examine predictability at the individual fund level.

To precisely assess the ability of each predictor j to forecast future fund returns, we decompose the fund population into three distinct categories:

•funds with unpredictable returns ($b_{i,j} = 0$);

•funds with predictable returns and a negative relation with predictor j ($b_{i,j} < 0$);

•funds with predictable returns and a positive relation with predictor j ($b_{i,j} > 0$).

Then, we measure the proportions of funds in the population, denoted by $\pi_R^0(j)$, $\pi_R^-(j)$, and $\pi_R^+(j)$, that fall into one of these three categories. The estimation procedure borrows from Barras, Scaillet, and Wermers (2010, henceforth BSW), and uses as input the estimated slope coefficients, $\widehat{b}_{i,j}$, across all funds. Importantly, this approach allows to measure true predictability, because it explicitly accounts for funds that exhibit predictability by luck alone (i.e., funds with statistically significant $\widehat{b}_{i,j}$, even when the true coefficient, $b_{i,j}$, equals zero). We refer the interested reader to BSW for further detail.

While the predictive regression in Equation (1) helps determine whether a given fund exhibits predictable returns, there are various sources for predictability. First, hedge fund benchmark expected returns (risk premia) can vary with changing economic conditions. Denoting by f_{t+1} the K-vector of portfolio-based benchmark excess returns in time $t+1$, we measure predictable risk premia using the regression

$$f_{t+1} = b_{f,0} + \sum_{j=1}^{J} b_{f,j} \cdot Z_{j,t} + u_{f,t+1}, \tag{2}$$

where $b_{f,0}$ is the K-vector of intercept coefficients, $b_{f,j}$ is the K-vector of slope coefficients associated with predictor j ($j = 1, ..., J$), and $u_{f,t+1}$ denotes the K-vector of factor innovations. While there is a large literature analyzing predictability of equity and bond factors (e.g., Fama and French (1989), Ilmanen (1995), and Ferson and Harvey (1999)), studying predictability of option-based factors included in hedge fund models is novel.

Second, hedge fund managers may have skills in security selection and benchmark timing that depend on the state of the economy. Indeed, Christopherson, Ferson, and Glassman (1998) and Avramov and Wermers (2006) document predictability of mutual fund managerial skills. If managers have specialized skills that best apply under specific economic conditions, their private information correlate with the predictive variables, making fund alphas predictable. To capture this intuition, we follow past work on mutual fund performance and model the dynamics of hedge fund return using

$$r_{i,t+1} = a_{i,0} + \sum_{j=1}^{J} a_{i,j} \cdot Z_{j,t} + \beta_i' f_{t+1} + \epsilon_{i,t+1}, \tag{3}$$

where $a_{i,0}$ is the intercept, $a_{i,j}$ is the alpha slope coefficient associated with each predictor, β_i the K-vector of fund risk loadings, and $\epsilon_{i,t+1}$ is the idiosyncratic fund-specific term. We decompose again the fund population into three predictability categories, now based on alpha variations, and denote by $\pi_\alpha^0(j)$, $\pi_\alpha^-(j)$, and $\pi_\alpha^+(j)$, the proportions of funds whose alphas are unrelated ($a_{i,j} = 0$), negatively related ($a_{i,j} < 0$), and positively related ($a_{i,j} > 0$) to predictor j, respectively.

To disentangle the two sources of hedge fund return predictability (benchmark returns versus alpha variation), we employ the restrictions imposed by the hedge fund benchmark model on the relation between the slope coefficients in Equations (1) and (3). Replacing f_{t+1} in Equation (3) with its expression in Equation (2), the predictive regression slope coefficient in Equation (1) becomes $b_{i,j} = a_{i,j} + \beta_i' b_j$, where $b_j = [b_{1,j}, ..., b_{K,j}]'$.

By comparing $b_{i,j}$ and $a_{i,j}$, we can easily determine the source of predictability for fund i. For one, if the explanatory power of predictor j is entirely driven by risk factors (as opposed to alpha), we would observe $b_{i,j} \neq 0$ and $a_{i,j} = 0$. This idea can be extended to examine the source of predictability in the entire cross-section of hedge funds by comparing the proportions of funds with predictable returns, $\pi_R^-(j)$ and $\pi_R^+(j)$, with the proportions of funds with predictable alphas, $\pi_\alpha^-(j)$ and $\pi_\alpha^+(j)$.

Given the large number of factors used in Equation (3) (typically seven factors in the Fung-Hsieh (2004) model), we assume that benchmark risk loadings are time-invariant. Using more parsimonious models, Bollen and Whaley (2009) and Patton and Ramadorai (2010) find that hedge fund betas are subject to structural breaks. Such breaks are less of a concern here since we are mostly interested in the estimated slope coefficients, $\widehat{a}_{i,j}$. While unmodeled beta variations can potentially bias the estimated unconditional alpha, they do not affect $\widehat{a}_{i,j}$ as long as the relation between the predictors and factors remains unchanged after the break.[2] To empirically verify this property, we examine the impact of changing betas associated with the prominent market and size factors. Following Fung et al. (2008), we allow for breaks after September 1998 and March 2000 and find in unreported results that the estimated proportion of predictable funds in the population remain virtually unchanged.

B Measuring the Economic Value of Predictability

B.1 The Trade-off between Unconditional and Predictable Performance

Previous studies on mutual funds and hedge funds typically rank and select funds based on expected fund performance. In a large population of hedge funds following different strategies, it is unlikely that funds with the highest unconditional mean are also those with the strongest predictability. In this multi-fund setting, investors willing to

[2] To see this, consider a simple model with one centered predictor, $z_{j,t} = Z_{j,t} - E(Z_{j,t})$, one factor, $f_{k,t+1}$, and one structural break at time $t+1 = \tau^*$: $r_{i,t+1} = \alpha_{i,0} + a_{i,j} \cdot z_{j,t} + \beta_{i,k} f_{k,t+1} + \beta^*_{i,k} f^*_{k,t+1} + \epsilon_{i,t+1}$, where $\alpha_{i,0}$ is the unconditional alpha and $f^*_{k,t+1} = f_{k,t+1} \cdot 1_{\{t+1 \geq \tau^*\}}$. Assuming that the relation between $z_{j,t}$ and $f_{k,t+1}$ is constant over time, we have $cov(z_{j,t}, f^*_{k,t+1} | f_{k,t+1}) = 0$, and the bias in $\widehat{a}_{i,j}$ in a constant-beta model is equal to zero: $E(\widehat{a}_{i,j} - a_{i,j}) = \beta^*_{i,k} \frac{cov(z_{j,t}, f^*_{k,t+1} | f_{k,t+1})}{var(z_{j,t} | f_{k,t+1})} = 0$. However, the estimated unconditional alpha is biased: $E(\widehat{\alpha}_{i,0} - \alpha_{i,0}) = \beta^*_{i,k} E(f^*_{k,t+1}) \neq 0$.

exploit predictability face a potential trade-off between unconditional and predictable performance. More formally, we can write the difference in expected returns between a conditional strategy with time-varying weights, $w^c_{i,t}$ $(i = 1, ..., M)$, and an unconditional strategy with constant weights, w^u_i, as

$$\mu^c - \mu^u = \sum_{i=1}^{M} cov(w^c_{i,t}, r_{i,t+1}) - \sum_{i=1}^{M} \left(w^u_i - E(w^c_{i,t})\right) \mu_i, \tag{4}$$

where μ_i is fund i unconditional (excess) mean. When investing in predictable funds, the investor tries to generate a positive covariance between the portfolio weights and the future returns of the predictable funds (the first term on the RHS). The second term on the RHS reflects the cost attributable to sacrificing unconditional performance of those funds that are excluded from the portfolio. Excessive tilts towards predictable funds could inflate this second term and make the conditional strategy unprofitable.

A simple way to incorporate this trade-off into the fund selection process is to rank funds according to the conditional (excess) mean, $\mu_{i,t} = E[r_{i,t+1}|Z_t]$, where Z_t stands for the J-vector of predictor values observed at the portfolio rebalancing time t. For simplicity, we start with the single-predictor case ($J = 1$), and discuss richer dynamics below. Denoting the centered predictor value by $z_{j,t}$ ($z_{j,t} = Z_{j,t} - E(Z_j)$, where $E(Z_j)$ is the predictor mean), we get

$$\mu_{i,t} = \mu_i + b_{i,j} z_{j,t}, \tag{5}$$

where μ_i is the fund unconditional mean and $b_{i,j} z_{j,t}$ is the predictable component. Equation (5) leads to a parsimonious use of predictive information, because predictable funds are only chosen when the predictive component, $z_{j,t} b_{i,j}$, is large enough, i.e., when $Z_{j,t}$ is sufficiently far away from $E(Z_j)$.

B.2 Implementing the Conditional Strategy

Implementing the conditional strategy involves three steps. The first step estimates, at each rebalancing time t and for each existing fund ($i = 1, ..., M_t$), the conditional mean, $\widehat{\mu}_{i,t}(j) = \widehat{\mu}_i + \widehat{b}_{i,j} z_{j,t}$, where the sample mean, $\overline{Z}_j$, replaces $E(Z_j)$ in the definition of $z_{j,t}$.

The conditional mean is not estimated with the same accuracy across funds with varying lives and investment strategies. To account for estimation uncertainty, the second step consists of computing the t-statistic of the estimated conditional mean:

$$t(\widehat{\mu}_{i,t}(j)) = \frac{\widehat{\mu}_{i,t}(j)}{\widehat{var}(\widehat{\mu}_{i,t}(j))^{\frac{1}{2}}}, \tag{6}$$

where $\widehat{var}(\widehat{\mu}_{i,t}(j))$ is the estimated variance of $\widehat{\mu}_{i,t}(j)$. The conditional t-statistic indicates how precisely the unconditional mean, μ_i, and the predictable component, $z_{j,t}b_{i,j}$ are estimated. Funds that exhibit higher $t(\widehat{\mu}_{i,t}(j))$ are likely to perform better.[3]

The third step of our dynamic setup consists of investing in the top decile of funds with the highest $t(\widehat{\mu}_{i,t}(j))$. This portfolio is held over the next period, after which the selection procedure is repeated (based on the new predictor value at time $t+1$). Our approach extends the decile portfolio approach of Elton, Gruber, and Blake (1996), Carhart (1997), which uses unconditional performance measures to rank funds.

Understanding the investment process is straightforward, once we decompose the conditional t-statistic. Specifically, let the unconditional t-statistic be $t(\widehat{\mu}_i) = \widehat{\mu}_i / \widehat{var}(\widehat{\mu}_i)^{\frac{1}{2}}$, and let the slope t-statistic be $t(\widehat{b}_{i,j}) = \widehat{b}_{i,j} / \widehat{var}(\widehat{b}_{i,j})^{\frac{1}{2}}$. Using Equation (6), we can write the conditional t-statistic of fund i as a weighted average of the unconditional and slope t-statistics:

$$\begin{aligned} t(\widehat{\mu}_{i,t}(j)) &= \left(\frac{\widehat{var}(\widehat{\mu}_i)}{\widehat{var}(\widehat{\mu}_{i,t}(j))}\right)^{\frac{1}{2}} t(\widehat{\mu}_i) + \left(\frac{\widehat{var}(\widehat{b}_{i,j})}{\widehat{var}(\widehat{\mu}_{i,t}(j))}\right)^{\frac{1}{2}} z_{j,t} \cdot t(\widehat{b}_{i,j}) \\ &= w_\mu(z_{j,t}) \cdot t(\widehat{\mu}_i) + w_{b,j}(z_{j,t}) \cdot t(\widehat{b}_{i,j}), \end{aligned} \tag{7}$$

where the weights, w_μ and $w_{b,j}$, depend on the difference, $z_{j,t}$, between the current value of the predictive variable, $Z_{j,t}$, and its long run mean $\overline{Z}_j$.

Essentially, when $Z_{j,t}$ is close to $\overline{Z}_j$, the strategy invests in the "unconditional" portfolio–the portfolio holding the top decile of funds with the highest unconditional t-statistic, $t(\widehat{\mu}_i)$. However, when $Z_{j,t}$ departs from $\overline{Z}_j$, the predictable component, $z_{j,t}b_{i,j}$, grows large, and the strategy invests in the "slope" portfolio–the portfolio holding the top decile of funds with the highest slope t-statistic, $t(\widehat{b}_{i,j}) \cdot sign(z_{j,t})$, where $sign(z_{j,t})$ denotes the sign of $z_{j,t}$. That is, the slope t-statistic is multiplied by the sign of $z_{j,t}$ to guarantee that the slope portfolio contains funds with the correct exposure ($\widehat{b}_{i,j} < 0$ when $z_{j,t} < 0$, and vice-versa).

In Figure 1, we confirm the existence of a trade-off between unconditional and pre-

[3] In an unconditional setting, the use of the t-statistic as an improved performance measure is advocated, among others, by Kosowski et al. (2006) and Kosowski, Naik, and Teo (2007).

dictable performance using our comprehensive hedge fund dataset (discussed below). To illustrate, Panel A plots the relation between the average t-statistic of funds included in the unconditional portfolio, denoted by $t(\widehat{\mu}_t^u)$, and the default spread (using other predictors provide similar insights).[4] To ease interpretation, $z_{j,t}$ is standardized, i.e., a value of one indicates that the predictor value is one standard deviation above its average. In a nutshell, we find that funds with high unconditional mean tend to be unpredictable. While the conditional t-statistic, $t(\widehat{\mu}_t^u)$, is high when $z_{j,t}$ is close to zero (driven by the high $t(\widehat{\mu}_i)$ across funds), it quickly goes down as $|z_{j,t}|$ increases because of the low and noisy estimate of the predictable component, $z_{j,t}\widehat{b}_{i,j}$ (i.e., $t(\widehat{b}_{i,j})$ is low across funds).

Panel B plots the relation between the conditional t-statistic of the slope portfolio, $t(\widehat{\mu}_t^s)$, and the default spread. In this case, we observe the exact opposite pattern: while $t(\widehat{\mu}_t^s)$ is low when $z_{j,t}$ is close to zero, it progressively increases as $|z_{j,t}|$ grows, implying that funds mostly likely to be predictable also exhibit low unconditional mean. Finally, Panel C illustrates the investment process described in Equation (7)–the conditional strategy moves away from the unconditional portfolio only when $|z_{j,t}|$ gets sufficiently large.[5]

Please insert Figure 1 here

B.3 The Combination Strategy: Dealing with Estimation Risk and Model Instability

The trade-off between unconditional and predictable performance suggests that investing in the slope portfolio is only advisable when $|z_{j,t}|$ is sufficiently large. But even in this situation, being fully invested in the slope portfolio may actually hurt performance for two reasons: estimation risk and model instability.

First, due to sampling errors it may be quite difficult to detect the (truly) predictable funds from the data (those funds with $z_{j,t}b_{i,j} > 0$). To illustrate, Panel C of Figure 1 clearly shows that the slope signal is less precise than the unconditional signal–the maximum value for the t-statistic of the unconditional portfolio equals 3.5, as opposed to only 1.6 for the slope portfolio. This suggests that when $|z_{j,t}|$ gets large, the conditional

[4] Specifically, for each fund i included in the unconditional portfolio in month t, we use past returns (over the last 36 months) to estimate $t(\widehat{\mu}_i)$, $\widehat{var}(\widehat{\mu}_i)$, $t(\widehat{b}_{i,j})$, and $\widehat{var}(\widehat{b}_{i,j})$. Then, we compute averages of these quantities across funds ($i \in$ unconditional portfolio in month t) and months ($t = 1, ..., T$). These averages are inserted in Equation (7) to compute $t(\widehat{\mu}_t^U)$, as a function of the predictor value, $z_{j,t}$.

[5] Note that for moderate predictor values $z_{j,t}$, some funds may have a high $t(\widehat{\mu}_{i,t}(j))$, despite having components, $t(\widehat{\mu}_i)$ and $t(\widehat{b}_{i,j})$ that, individually, do not belong to their respective top deciles. While these funds are chosen by the conditional strategy, they do not belong to the unconditional and slope portfolios.

strategy may simply trade funds with (truly) high unconditional performance for funds exhibiting low levels of predictability.

Second, even if at some point in time, we can precisely estimate the predictive regression slope coefficients, numerous factors, such as the investors' search for successful forecasting models, technological shocks, or institutional changes, make the predictive model unstable (e.g., Timmermann (2008)). Since we expect both the identity of the relevant predictors as well as their associated slope coefficients to change over time, the single-predictor model considered so far may not capture all variation due to changing economic conditions. If the forecasting power of a given predictor follows short-term cyclical patterns, there are periods when the estimated predictable component, $z_{j,t}\widehat{b}_{i,j}$, conveys wrong signals about future performance, leading to a poor fund selection.

To address estimation risk, model instability, and the tradeoff between unconditional and predictable performance, we implement an alternative conditional strategy building on the combination forecast literature (e.g., Bates and Granger (1969), Timmermann (2006)). Specifically, for each existing fund i at the rebalancing time t ($i = 1, ..., M_t$), we estimate its conditional t-statistic, $t(\widehat{\mu}_{i,t}(j))$, using each predictor j separately ($j = 1, ..., J$). Second, we compute the simple average across all J conditional t-statistics:

$$t(\widehat{\mu}_{i,t}) = \frac{1}{J}\sum_{j=1}^{J} t(\widehat{\mu}_{i,t}\,(j)) = w_{\mu} \cdot t\,(\widehat{\mu}_i) + \frac{1}{J}\sum_{j=1}^{J} w_{b,j}(z_{j,t}) \cdot t(\widehat{b}_{i,j}), \tag{8}$$

where $w_{\mu} = \frac{1}{J}\sum_{j=1}^{J} w_{\mu}(z_{j,t})$.[6] We ultimately design an investment strategy that holds the top decile of funds with the highest combination t-statistic, $t(\widehat{\mu}_{i,t})$.

This *combination* strategy exhibits several appealing properties. First, since it is unlikely to observe extreme values for all predictors simultaneously, the total weight, w_{μ}, associated with the fund unconditional t-statistic, $t\,(\widehat{\mu}_i)$, remains high. At the same time, the importance of the slope signals in the investment process decreases because each individual weight, $w_{b,j}(z_{j,t})$, is divided by J. The combination strategy shrinks the portfolio towards the unconditional portfolio and reduces the impact of both estimation risk and model instability (see Rapach, Strauss, and Zhou (2010) for a discussion in a single-asset setting). Second, similar to portfolio diversification, combining forecasts generally reduces the out-of-sample forecast error variance (see Timmermann (2006)). In particular, Hendry and Clements (2004) show that combining forecasts provides a good

[6] While more complex weighting schemes exist, the simple average tends to perform well, as the weights do not have to be estimated (see Timmermann (2006)). As an alternative to combination forecast, previous papers use Bayesian averaging, where the weight associated with each predictive model depends on the model prior distribution (e.g., Avramov (2002)).

hedge against structural breaks in the data generating process. Hence, the combination strategy should be more likely to consistently detect those funds with predictable returns.

Alternatively, we could try to improve on the single-predictor model by including all predictors simultaneously in the predictive regression. However, as shown by Avramov (2002) and Goyal and Welch (2008), the out-of-sample forecast errors of this multi-predictor model are large, as multiple slope coefficients are estimated with less accuracy. Since there is no shrinkage towards the unconditional portfolio, performance deteriorates because large positions are taken in slope portfolios that exhibit low or no predictability out-of-sample.[7] However, for comparative purposes, we also report the performance of the conditional strategy based on the multi-predictor model.

C Estimation Issues

C.1 Correcting for Small Sample Bias

It is well-known that the ordinary least-square (OLS) estimation of the predictive regression slope coefficients is subject to small-sample bias. That is, the expected value of estimated slope coefficients, $E(\widehat{b}_i^{ols}) = E(\widehat{b}_{i,1}^{ols}, ..., \widehat{b}_{i,J}^{ols})'$, is different from its true parameter value, b_i (e.g., Stambaugh (1999)). While this bias disappears in large enough samples, it is an important concern here because the return history for many hedge funds is short. For one, survivorship bias-free databases only start in January 1994.

To illustrate, we estimate that for 25% of the funds in the population, a one-standard deviation increase in the dividend yield, leads to an upward bias in expected return greater than 3.1% on an annual basis. This bias is even more pronounced for specific investment categories. Ignoring small sample bias is likely to affect the estimated proportions of predictable funds, as well as the fund estimated t-statistics, leading to a poor fund selection.

The small sample bias arises under two conditions frequently met empirically: 1) the predictors are persistent, e.g., Z_{t+1} has an autoregressive VAR(1) structure:

$$Z_{t+1} = \theta + \Phi Z_t + \upsilon_{t+1}, \tag{9}$$

where Φ is the $J \times J$ companion matrix, and υ_{t+1} is the J-vector of innovation; 2) the hedge fund innovation, $u_{i,t+1}$, is contemporaneously correlated with υ_{t+1}. That is, we

[7] To see why there is no shrinkage, note that in the multi-predictor case, the conditional t-statistic has a similar expression as in Equation (7): $t(\widehat{\mu}_i) = w_\mu(z_t) \cdot t(\widehat{\mu}_i) + w_b'(z_t) \cdot t(\widehat{b}_i)$, where z_t, w_b, and $t(\widehat{b}_i)$ are all J-vectors. When the k^{th} element of z_t gets large, the conditional strategy invests in the k^{th} slope portfolio that holds funds with the highest k^{th} slope t-statistic, $t(\widehat{b}_{i,k})$.

can express the hedge fund innovation using $u_{i,t+1} = \phi_i' \upsilon_{t+1} + e_{i,t+1}$, where ϕ_i denotes the J-vector of innovation coefficients, and $e_{i,t+1}$ is the fund residual term (orthogonal to both υ_{t+1} and Z_t).[8] Since the OLS-estimated companion matrix, $\widehat{\Phi}$, is biased in small samples (Nicholls and Pope (1988)), the slope estimate, $\widehat{b}_i$, inherits some of the bias in $\widehat{\Phi}$ because of condition 2). Following Stambaugh (1999), this bias can be written as[9]

$$bias(\widehat{b}_i^{ols}) = E(\widehat{b}_i^{ols} - b_i) = E\left(\widehat{\Phi} - \Phi\right)' \phi_i. \tag{10}$$

Intuitively, Equation (10) can be interpreted as an omitted variable bias, since $\widehat{b}_i^{ols}$ captures the influence of the omitted innovation vector, $\upsilon_{t+1,}$ on $r_{i,t+1}$. Therefore, as noted by Amihud and Hurvich (2004) and Amihud, Hurvich, and Wang (2008, AHW hereafter), if we include the J-vector υ_{t+1} as an additional explanatory variable and write

$$r_{i,t+1} = b_{i,0} + \sum_{j=1}^{J} b_{i,j} Z_{j,t} \ + \phi_i^{'} \upsilon_{t+1} + e_{i,t+1}, \tag{11}$$

the small-sample bias disappears as the orthogonality holds, i.e., $E\left(e_i \left| X, V\right.\right) = 0$, where $e_i = [e_{i,2}, ..., e_{i,T+1}]'$, $X = [(1, Z_1', ..., Z_T']$, and $V = [\upsilon_2', ..., \upsilon_{T+1}']$. Of course, we cannot observe the true innovation vector, thus we compute a proxy, υ_{t+1}^c, using the procedure proposed by AHW. After replacing υ_{t+1} with υ_{t+1}^c, we can compute the bias-corrected estimated slope coefficients, $\widehat{b}_{i,j}$, by applying standard OLS to the augmented regression in Equation (11). Using extensive simulation tests, AHW find that their approach achieves a substantial reduction in the small-sample bias ($\widehat{b}_{i,j}$ is not totally bias-free though, as we use υ_{t+1}^c instead of the true υ_{t+1}). While other approaches are also feasible, such as the bootstrap (Nelson and Kim (1993)), the AHW procedure is computationally much faster as it boils down to estimating a single regression for each fund. Given the great number of funds in our sample, computational efficiency is strongly appealing.

The framework proposed by AHW focuses on traditional predictive regressions. Here, we extend their methodology on several fronts: 1) to examine predictability in a richer setting that incorporates alpha predictability and different time horizons (monthly and

[8] Unless the J-vector of predictor innovations, υ_{t+1}, is strongly correlated with news about current and future expected cash flows, there must be a contemporaneous correlation between υ_{t+1} and $u_{i,t+1}$. Changes in future expected returns captured by υ_{t+1} affect both prices and the contemporaneous fund return, $r_{i,t+1}$ (e.g., Cochrane (2008) and Pastor and Stambaugh (2009)).

[9] Using $X_{(T\times J+1)} = [(1, Z_1')', ...(1, Z_T')']'$, and $Y_{i(T\times 1)} = [r_{i,2}, ..., r_{i,T+1}]'$, we can write $Y_i = Xb_i + V\phi_i + e_i$, where $V_{(T\times J)} = [\upsilon_2', ..., \upsilon_{T+1}']'$. Replacing V with $Z - X\Phi'$, where $Z_{(T\times J)} = [Z_2', ...Z_{T+1}']'$, we can use the standard OLS formula to get $E(\widehat{b}_i^{ols}) = E\left((X'X)^{-1}X'Y_i\right) = b_i + E\left((X'X)^{-1}X'\left(Z - X\Phi'\right)\phi_i\right) = b_i + E\left(\widehat{\Phi} - \Phi\right)' \phi_i$.

quarterly); 2) to account for potential hedge fund illiquidity (discussed below). All the technical details on estimation (i.e., the AHW approach and related extensions) are detailed in Appendix A.

C.2 Accounting for Hedge Fund Illiquidity

Some hedge funds invest in illiquid assets, such as emerging market debt, asset-backed securities, or over-the-counter derivatives. Such assets may be affected by non-synchronous trading (stale prices) and may also facilitate return misreporting activities, as documented by Bollen and Pool (2009). Illiquidity tends to smooth hedge fund returns over time, and induce positively correlated residuals, $e_{i,t+1}$, in Equation (11) (see Getmansky, Lo, and Makarov (2004)). All else equal, the standard deviation of the estimated coefficients, $\widehat{b}_{i,0}$, and $\widehat{b}_{i,j}$, is higher for funds with positively correlated residuals. Failing to adjust for this correlation, we may wrongly conclude, based on the t-statistics of $\widehat{b}_{i,0}$, and $\widehat{b}_{i,j}$, that illiquid funds generate a higher unconditional mean and/or have highly predictable returns.

To explicitly control for illiquidity when computing the variance of $\widehat{b}_{i,0}$, and $\widehat{b}_{i,j}$, we model the residual, $e_{i,t+1}$, of each fund i as an $AR(p)$ process ($i = 1, ..., M$):

$$e_{i,t+1} = \sum_{l=1}^{p} \rho_{i,l} e_{i,t+1-l} + \xi_{i,t+1}, \tag{12}$$

where $\rho_{i,l}$ is the autoregressive coefficient at lag l ($l + 1, ..., p$), and $\xi_{i,t+1}$ is the innovation term. Using the estimated residuals, $\widehat{e}_{i,t+1}$, to obtain consistent autoregressive coefficient estimates, we follow the BSW approach to compute the proportions of funds with non-zero coefficients ($\rho_{i,l} \neq 0$) at different lags. In unreported results, we find that 29.1% and 28.1% of the funds have a one-month and two-month lag coefficients different from zero, respectively, while the proportion falls to 4.9% at a three-month lag. The results are qualitatively similar across investment styles (some categories such as convertible arbitrage exhibit higher proportions, consistent with Getmansky, Lo, and Makarov (2004)). Based on this evidence, we use an AR(2) model that we estimate for each fund separately to control for the cross-fund coefficient variation observed in the data. As an alternative to the AR specification, we also compute the variance of the estimated regression coefficients using the Newey-West (1987) methodology, and find that the results (to be presented) remain unchanged.

III Data Description

We use four economically motivated instruments to predict future hedge fund returns: the default spread, the dividend yield, the VIX, and aggregate fund flows. Given the relatively small number of monthly observations for hedge funds, model parsimony is an important consideration in our choice of predictors. Parsimony also avoids the search over a large number of predictors which could invoke data-mining concerns.[10] All predictors are observed at a monthly frequency and appropriately lagged to forecast future hedge fund excess returns.

The default spread is the yield differential between Moodys BAA- and AAA-rated bonds. Previous studies (e.g., Keim and Stambaugh (1986)) show that the default spread can predict future stock and bond returns. The dividend yield is the total of annual cash dividends on the value-weighted CRSP index divided by the current index level. Fama and French (1989) suggest that the dividend yield is a business cycle indicator that peaks in recession when expected returns required by investors are high. We also use the VIX index from the CBOE. Taylor, Yadav, and Zhang (2010) present evidence that implied volatilities help predict stock returns. Moreover, volatility may capture some of the option-like features in hedge fund returns (e.g., Agarwal and Naik (2004)). Aggregate flows are calculated as the value-weighted percentage net inflows into our sample of hedge funds. As discussed in Naik, Ramadorai, and Stromqvist (2007) and Fung et al. (2008), new money can create capacity constraints, leading to lower future returns.

Figure 2 shows that during the financial crisis of 2008, the dividend yield, the default spread, and the VIX exhibited extreme deviations from their past historical average. For this reason, we initially focus on the period 1994-2007 to assess hedge fund predictability, and run a separate analysis to check the robustness of our results during the 2008 crisis.

Please insert Figure 2 here

Panel A of Table I reports descriptive statistics for the hedge funds included in our sample during the January 1994 through December 2007 period. We evaluate hedge fund performance based on monthly net-of-fee returns using a new database that aggregates, for the first time, data reported by five different providers (BarclayHedge, TASS, HFR, CISDM and MSCI). To create this data set, we carefully control for a number of potential

[10] Patton and Ramadorai (2010) address the need for a parsimonious model by examining a large number of predictors and risk factors, and, for each fund, selecting one predictor and two factors based on a statistical criterion.

biases discussed in Appendix B. To estimate the predictive regression coefficients, each fund is required to have at least 36 monthly return observations, leading to a final sample of 7,991 funds.[11]

We observe from Panel B that the default spread, the dividend yield, and the VIX exhibit high positive autocorrelation ($\rho = 0.95$, 0.97, and 0.84, respectively). These large coefficients highlight the importance of controlling for small sample bias in the estimation process. In addition, correlations across predictors are low on average, suggesting that each variable captures specific variations in economic conditions. In this context, the combination strategy could add value over individual predictors.

Finally, Panel C contains summary statistics for the benchmarks included in the Fung and Hsieh (2004) seven factor model. Equity Market is the S&P 500 return minus risk free rate, Equity Size is the Wilshire small cap minus large cap return, Bond Term is the change in the constant maturity yield of the 10–year Treasury appropriately adjusted for duration (to represent returns on a traded portfolio), Bond Default is the change in the spread of Moody's Baa minus the 10–year Treasury (also adjusted for duration), and Trend Bond, Currency, and Commodity are the straddle-type trend following strategies.[12]

Please insert Table I here

IV Empirical Results

A Predicting Hedge Fund Performance

A.1 Return versus Alpha Predictability

We begin our empirical analysis by measuring in-sample return and alpha predictability across individual hedge funds over the period 1994-2007. While Panel A of Table II reports the evidence for all funds in the population, Panels B to K focus on the different investment categories. For each panel, the first row (Return) contains the estimated proportions of funds with predictable returns, $\pi_R^-(j)$ and $\pi_R^+(j)$, associated with each predictor using the (bias-corrected) estimated slope coefficients, $\widehat{b}_{i,j}$, in Equation (1). The last row-element displays the proportion of predictable funds using all predictors simultaneously, $\widehat{\pi}_R^{Joint}$. Similarly, the second row of each panel (Alpha), reports the estimated proportions of funds with predictable alphas, $\pi_\alpha^-(j)$, $\widehat{\pi}_\alpha^+(j)$, and $\widehat{\pi}_\alpha^{Joint}$, obtained from the (bias-corrected) estimated slope coefficients, $\widehat{a}_{i,j}$, in Equation (3).

[11] While this requirement may lead to survivorship bias, unreported results show that our results are robust to using funds with 24 monthly observations.

[12] We thank David Hsieh for making these factors available on his website.

Overall, two insights stand out from Table II. First, there is ample evidence of predictability. In the entire population about 60.5% of the funds are predictable, while this proportion ranges from 31.1% (managed futures) to 83.7% (convertible arbitrage) across investment styles. Second, predictability is primarily attributable to alpha variation. Comparing the Return and Alpha rows reveals that for most hedge fund styles, the proportions of funds exhibiting return and alpha predictability are almost identical (i.e., $\widehat{\pi}_R^- \approx \widehat{\pi}_\alpha^-$ and $\widehat{\pi}_R^+ \approx \widehat{\pi}_\alpha^+$). This interpretation is corroborated by evidence that the Fung-Hsieh benchmark factors are largely unpredictable. Unreported results show that only 5 out of the 28 estimated slope coefficients in Equation (2) are significant (at the 10% level). Sensitivity tests discussed below show that alpha predictability is robust to augmenting the Fung-Hsieh model with additional risk factors.

Please insert Table II here

Several investment styles exhibit a high fraction of funds that are positively predictable by the default spread. For instance, emerging market funds display the strongest relation with $\widehat{\pi}_R^+ = 33.5\%$. Focusing on emerging market funds, the median of the (standardized) slope coefficient, $\bar{b}_j$, implies that a one standard deviation increase in the default spread causes an additional 43 bp per month (5% per year) investment payoff. Widening credit spreads typically coincide with flight to quality, which could in turn forecast higher future returns. A similar reasoning holds for the carry trade strategies in FX markets followed by global macro funds ($\widehat{\pi}_R^+ = 13.2\%$). Widening credit spreads can trigger the unwinding of carry trades, which leads to increasing future expected returns (e.g., Jylha and Suominen (2010)).

The dividend yield shows a very consistent pattern across all styles, except for managed futures. While only a few funds have positive slope coefficients, the majority displays a negative exposure ($\widehat{\pi}_R^-$ ranges from 14.8% to 37.8%). One explanation, consistent with the business cycle interpretation of the dividend yield, is the role of leverage in explaining hedge fund performance. In recessions, when the dividend yield is high, leverage availability from prime brokers is constrained, forcing hedge funds to reduce their exposures, thus generating low returns and alphas (see Ang, Gorovyy, and van Inwegen (2010)).[13] This could also explain the opposite pattern for managed futures funds. Such funds encounter fewer leverage constraints as they trade in futures markets that are reasonably liquid over the business cycle. Liquid futures positions allow

[13] Our results contrast with the positive relation between the dividend yield and future mutual fund returns (Ferson and Qian (2004)). Mutual funds have a strong exposure to the stock market, whose returns tend to be positively related to the dividend yield (e.g., Fama and French (1989)).

managed futures funds to easily adjust their desired margin to equity ratio.

The vast majority of fund styles exhibit a negative relation with the VIX. For instance, 39.0% of event-driven funds are negatively associated with the VIX, consistent with the higher deal failure rate in turbulent periods (e.g., Mitchel and Pulvino (2001)). Convertible arbitrage funds also display negative exposure ($\widehat{\pi}_R^- = 24.9\%$) as an increasing VIX may reduce opportunities of cheap volatility. Indeed, as shown by Choi et al. (2010), convertible arbitrage funds often exploit mispriced volatility in convertible bonds. Managed futures funds, in contrast, benefit from higher uncertainty ($\widehat{\pi}_R^+ = 16.7\%$). Indeed, managed futures funds (which include CTAs) tend do well following trend reversals which occur during periods of extreme VIX values. To illustrate, during the 2008 financial crisis, a Financial Times article observed that "*CTAs as an investment typically do best in periods of market chaos[...] they performed relatively well when implied volatility of equities rose on fears of Russian default and the Long Term Capital Management crash in 1998. A similar picture emerged in the run up to the Iraq war in 2002-2003, and now the credit crunch of 2008*".[14]

The strongest evidence for negative return predictability is found for aggregate flows. Looking at the entire population, an increase in flows leads to a decrease in returns and alphas for 33.2% and 26.7% of the funds, respectively. This flow-return relation confirms the finding of Naik, Ramadorai, and Stromquist (2007) who argue that capacity constraints, attributable to excessive inflows, hurt performance.[15] Unsurprisingly, this effect is particularly strong for the crowded market of convertible arbitrage that is dominated by hedge funds ($\widehat{\pi}_R^- = 44.7\%$). It is also strong for funds of funds, suggesting that they struggle to generate performance by deploying capital after large fund inflows.

One important insight coming out from Table II is the prominent asymmetry in the direction of predictability. For instance, most predictable funds exhibit a negative slope coefficient with respect to the dividend yield. Thus, if the dividend yield is well above its long run mean, it will be difficult for the conditional strategy to detect funds with the correct exposure (i.e., funds with $b_{i,j} > 0$). The same pattern holds for the other predictors. We will show later that this asymmetry carries strong implications for understanding out-of-sample predictability.

Arguably, liquidity constraints (such as lock-up periods) may prevent investors from rebalancing their fund portfolio at a monthly frequency.[16] For hedge fund predictability

[14] "One investment that loves chaos", Financial Times, 23 November 2008.

[15] The existence of capacity constraints is confirmed by our subperiod analysis: while $\widehat{\pi}_\alpha^-$ equals 6.5% of the entire population between 1994 and 2000, it jumps to 38.9% during the most recent period (2001-2007).

[16] Since liquidity constraints prevent investors from investing additional money into hedge funds, it

to be exploitable out-of-sample, it must, therefore, be present at lower frequencies. As a first test, we examine predictability at a quarterly frequency. We find in unreported results that quarterly returns are still predictable, which could support the economic value of hedge fund return predictability. This is what we examine next.

B The Economic Value of Predictability

B.1 Performance Analysis

While we document ample evidence of in-sample predictability, institutional investors (such as funds of funds) typically encounter nontrivial impediments to exploiting hedge fund return predictability. First, predictable funds may have low unconditional mean, leading to a trade-off between unconditional and conditional performance. Second, the identity of predictable funds is unknown, and must be inferred from the data (estimation risk). Third, in-sample predictability may not carry forward if the predictive relation changes over time (model instability). To address these concerns, we carry out a range of out-of-sample tests that carefully incorporate real-world investment constraints and quantify the economic value of predictability.

First, we account for liquidity constraints by excluding closed funds as well as considering a one-year lock-up period (i.e., annual rebalancing). Second, there is a practical limit to the number of individual funds held by funds of funds. While data on such holdings is not publicly available, Lhabitant (2006) indicates that the typical number is about 40. As a result, we limit the minimum and maximum number of funds in the portfolio to 25 and 75, respectively. Third, investors typically do not invest in funds that are too small relative to their own size. As discussed in Ganshaw (2010), few institutional investors want to represent more than 10% of a fund's assets under management. We use this rule to set up a dynamic AuM cutoff equal to the minimum fund size such that a "typical" fund of funds does not breach this 10%-threshold.[17] The resulting cutoff rises from $12 million at the beginning to $63 million at the end of our sample. Contrary to the constant cutoffs used in previous studies (e.g., $20 million), our filter explicitly accounts for the growth in the hedge fund industry over time. Finally, we exclude funds of funds since institutional investors often focus on individual funds to avoid extra fees.

We consider the following trading strategies: (i) the unconditional portfolio which

may also explain why alphas are predictable in the first place. Baquero and Verbeek (2009) and Ding et al. (2009) raise a similar argument when they examine the performance of "smart money" strategies in hedge funds.

[17] The "typical" fund has an average size (as measured from the funds of funds AuM in our sample) and invests in 50 funds (the midpoint in our investment strategy).

uses the unconditional t-statistic; (ii) the single-predictor strategy which ranks funds according to the conditional t-statistic obtained from each predictor (default spread, dividend yield, VIX, and aggregate flows); (iii) the multi-predictor strategy which uses all predictors simultaneously; (iv) the combination strategy which averages across the single-predictor t-statistics. The construction of the different portfolios proceeds as follows. At the end of each year, we estimate, for each strategy, the appropriate signal for each existing fund using past three-year returns, and form the top decile portfolio. This portfolio is kept during one year, after which the entire procedure is repeated.

Panel A of Table III reports the out-of-sample performance of all investment strategies between January 1997 and December 2007 (the period 1994-1996 is used for the initial estimation). First, we observe that the unconditional portfolio generates solid performance–the annual Fung-Hsieh alpha and Information Ratio (IR) are 5.8% and 2.4, respectively, thus easily beating the value- (VW) and equal-weighted (EW) hedge fund indices. This performance reflects the high signal accuracy of the unconditional mean (see Panel C of Figure 1), and is consistent with the results obtained by Kosowski, Naik, and Teo (2007), and Jagannathan, Malakhov, and Novikov (2010). Second, none of the single-predictor strategies outperforms the unconditional portfolio on a risk-adjusted basis (Information and Sharpe Ratios). Third, consistent with the previous literature (e.g., Avramov (2002) and Goyal and Welch (2008)), the multi-predictor strategy produces the worst performance ($IR = 1.5$) most likely due to severe estimation errors.

Remarkably, the combination strategy performs well. For one, Information and Sharpe Ratios, which amount to 2.7 and 1.9, respectively, are statistically significantly higher than their unconditional counterparts (at the 5% level).

Please insert Table III here

This superior performance translates into large economic gains over the entire period as shown in Figure 3. We show that the wealth produced by the combination strategy accumulates steadily over time and reaches \$3.74 (for each dollar invested in 1997), leading to an additional 50 cents of terminal wealth compared to the unconditional portfolio.

Please insert Figure 3 here

Repeating this analysis across the two largest investment categories (long-short equity and directional funds, which includes global macro, managed futures, and emerging markets), we confirm, in unreported results, the superior performance of the combination strategy.

Importantly, the combination strategy does not introduce higher tail risk. The 1% and 5% Value-at-Risk in Table III equals -1.4% and -0.7% per month, respectively, and are thus the lowest among all competing strategies. From an investor perspective, the superior performance of the combination strategy is accompanied by several additional advantages displayed in Panel B. First, the strategy does not involve extensive (and possibly unrealistic) portfolio turnover–it exhibits the second lowest (66%) annual turnover of constituent funds. Second, the autocorrelation coefficients indicate that the superior performance of the combination strategy is not due to holding illiquid funds. The coefficients are comparable in magnitude to the other strategies, and are lower than the typical hedge fund autocorrelation coefficients (e.g., Lo (2002)). Finally, the performance of the combination strategy is not driven by concentrated bets on specific investment categories. The highest weight, invested in long-short equity funds, is only equal to 11.8% on average over the period.

In a recent study, Avramov, Kosowski, Naik, and Teo (2010; AKNT) form hedge fund portfolios that exploit predictability using an optimizing Bayesian framework. A comparison of the combination strategy with the Bayesian methodology reveals important differences. For one, the combination strategy invests in a larger number of funds (68 against 12 in AKNT), suggesting that it may be more robust to extreme market conditions, as discussed in Jagananthan and Ma (2003). Furthermore, Panel C of Table III shows that the combination strategy behaves more like a pure alpha strategy, as its exposures to the Fung-Hsieh risk factors are extremely low. More importantly, the investment strategies considered here are very intuitive, as the investment mechanism boils down to considering two portfolios only, the unconditional and slope portfolios. Thus, our proposed setup allows for a deeper understanding of the drivers of out-of-sample performance and the links between in- and out-of-sample forecast quality, as shown below.

B.2 Explaining the Performance of Single-Predictor Strategies

When the predictor value departs from its long run mean, the single-predictor strategy invests in the slope portfolio that includes funds with high estimated predictable components, $\widehat{b}_{i,j} \cdot sign(z_{j,t})$. Therefore, a key driver of performance should be the slope portfolio's ability to detect the (truly) predictable funds from the data.

Consider, for instance, the default spread. In Table IV, we group the values of the default spread, $z_{j,t}$, observed at each rebalancing date into quintiles, and examine the strategy allocation when $z_{j,t}$ is either far below (bottom quintile) or above (top quintile) its long run mean. As expected, Panel A shows that when the default spread takes on extreme values, the conditional strategy has a large exposure to the slope portfolio. For instance, when $z_{j,t}$ is in the top quintile, 77.2% of the funds chosen by the strategy are common to those in the slope portfolio ($\%_{slope} = 77.2\%$). But more importantly, the quality of the predictive information is very different depending on the sign of $z_{j,t}$. Looking at the cross-fund average slope t-statistic, denoted by $t(\widehat{b}_j)$, we observe that it is quite low when $z_{j,t}$ is negative ($t(\widehat{b}_j) = 0.96$), and very high when $z_{j,t}$ is positive ($t(\widehat{b}_j) = 2.03$). This reflects the strong asymmetry documented in-sample. Since there are a lot more funds with a positive exposure to the default spread in the population (in Panel A of Table II, $\pi_R^+(j) = 14.5\%$ versus $\pi_R^-(j) = 1.6\%$), it should be much easier to identify these predictable funds from the data when $z_{j,t}$ is positive.

As suggested, Panel B reveals that the performance of the default spread strategy is strongly dependent on the signal accuracy of the slope portfolio. The Fung-Hsieh annual alpha, for instance, is equal to 14.8% after observing a high default spread at the rebalancing date (as opposed to 6.0% when it is low).

The same analysis applies to the dividend yield, VIX, and aggregate flows. In each case, Panel A documents a strong asymmetry in the quality of predictive information, which matches the in-sample asymmetry observed in Table II. This, in turn, leads to a strong asymmetry in future performance, as shown in Panel B. Observing such a clear-cut pattern in performance is striking given that all measures are computed out-of-sample.

Please insert Table IV here

Figure 4 provides additional evidence on the asymmetric relation between the predictor value and out-of-sample performance. In each Panel, we plot the evolution of the predictor value along with the difference in Information Ratio between the single-predictor strategy and the unconditional portfolio. The results confirm those in Table IV. For instance, the performance of the dividend yield strategy goes up in times when the dividend yield is below average, consistent with the negative exposure to the dividend yield documented in-sample ($\pi_R^-(j) = 22.6\%$, as opposed to $\pi_R^+(j) = 3.6\%$).

Please insert Figure 4 here

B.3 Explaining the Performance of the Combination Strategy

In essence, taking large positions in the slope portfolio hurts performance when the quality of the predictive information is poor. The combination strategy overcomes this difficulty by taking a substantial position in the unconditional portfolio. To illustrate,

Figure 5 shows the fraction of funds chosen by each conditional strategy that are common to those held by the unconditional portfolio. While there are large variations for both single- and multi-predictor strategies, the proportion associated with the combination strategy is stable around 80%.

While Rapach, Strauss, and Zhou (2010) highlight the benefits of shrinkage in a single-asset environment, we argue that the combination approach is even more attractive in the presence of multiple assets because of the trade-off between unconditional and predictable performance. Indeed, by investing in the slope portfolio at the wrong time, the investor picks up funds with low predictability, and, in addition, fails to capture the relatively high performance produced by the unconditional portfolio.

Please insert Figure 5 here

How valuable are the active bets taken by the active portfolio? As any active strategy, the return of the combination strategy, r_{t+1}^{com}, equals the return of the unconditional portfolio, r_{t+1}^{u}, plus the return of an investment in a zero-cost long-short portfolio, r_{t+1}^{ls}:

$$r_{t+1}^{com} = r_{t+1}^{u} + w_t \cdot r_{t+1}^{ls} = r_{t+1}^{u} + w_t \cdot (r_{t+1}^{l} - r_{t+1}^{s}), \tag{13}$$

where r_{t+1}^{l} is the return of a long portfolio that contains those funds selected by the combination strategy, but not by the unconditional portfolio, while r_{t+1}^{s} is the return of a short portfolio consisting of funds that are only selected by the unconditional strategy (but not by the combination portfolio). Finally, w_t is the weight invested in this long-short position. Observe from Figure 5 that w_t is around 20%. Intuitively, the construction of the long and short portfolios can be described as a two-step process. First, given its higher focus on unconditional performance, the combination strategy preselects funds having a sufficiently high unconditional signal, $t\left(\widehat{\mu}_i\right)$. Then, in this subset, it looks for predictable funds using all predictors. Funds having a positive (and precisely measured) predictable component are included in the long portfolio, while funds with low predictable component are put in the short portfolio.

Is performance of the combination strategy subject to the asymmetry described earlier? We rule out this possibility. Specifically, for each predictor, Table V displays the annual out-of-sample performance of the long and short portfolios, r_t^l and r_t^s, when the predictor value, $z_{j,t}$, is far below (bottom quintile) or above (top quintile) its long run mean. For instance, we find that the Fung-Hsieh alpha differential against the unconditional portfolio, $\widehat{\alpha} - \widehat{\alpha}_U$, is equal to 5.3% per year, even when the default spread has the "wrong" value, i.e., when it is below average ($z_{j,t} < 0$). By combining across predictors, the long portfolio is able to generate a high performance that is orthogonal to that of the slope portfolios. Indeed, Panel A shows that the proportion of funds that are common to any given slope portfolio is at most equal to 14%. By the same token, Panel B shows that the combination strategy also excludes funds that are expected to underperform.[18]

Please insert Table V here

C Impact of the 2008 Financial Crisis

The 2008 financial crisis is arguably the biggest crisis in modern financial history. During this period, a dramatic deviation of all predictors from their historical means had been recorded (see Figure 2). Because of these strong predictive signals, the 2008 crisis provides us with a natural experiment to test the usefulness of conditioning information. For investors, it is of great practical importance to answer at least two questions. First, is the superior performance of the combination strategy robust to the inclusion of 2008? Second, which of the conditional strategies, if any, would have anticipated the 2008 events?

Panel A of Table VI reports the out-of-sample performance of the unconditional and conditional strategies (single-, multi-predictor, and combination) between January 1997 and December 2008. Even including the crisis, the combination strategy still achieves a reasonably strong performance. Its Information Ratio is not only the highest (IR =1.9), but is also statistically significantly higher than that of the unconditional portfolio (the associated p-value is below 1%). In general, all strategies have been hit quite hard during the crisis. For instance, the cumulative loss during the final quarter of 2008 amounts to 14.1% for the unconditional portfolio, and to 12.3% and 14.1% for the default spread and aggregate flows strategies, respectively. There is one noticeable exception. The conditional strategy based on the VIX yields the highest Sharpe Ratio (SR=1.4), and resists remarkably well during the final quarter of 2008 (with only a 3.1% cumulative loss).

Please insert Table VI here

It is interesting to understand the performance drivers of the VIX strategy. To this end, we compare the characteristics of the single-predictor strategies in 2008. We observe in Panel B of Table VI that the December 2007 value of all predictors was very

[18] As an additional test of the combination strategy's ability to detect predictable funds, we examine the out-of-sample forecasting ability of the monthly long-short portfolio returns, r^{ls}_{t+1}. These unreported results confirm that combined forecasts, averaged across the four predictors, achieve a lower Mean Squared Prediction Error (MSPE) than the historical mean.

different from their average, leading to a large investment in the slope portfolios (the proportion of common funds, $\%_{slope}$, ranges from 48.0% to 76% across predictors). How appropriate were these slope portfolios in 2008? To answer this question, we compute the hit ratio that determines the proportion of months in 2008 when the portfolio predictable component, $\widehat{b}_{j,t} \cdot z_{j,t}$, is positive, where $z_{j,t}$ is the predictor value at the start of each month, and $\widehat{b}_{j,t}$ is the average slope coefficient among the funds included in the slope portfolio ($\widehat{b}_{j,t}$ is estimated with the 36 most recent monthly observations). The results suggest that the VIX strategy performs well in 2008 because its predictive power is much more stable–its hit ratio is extremely high (83.3%), whereas the hit ratios of the default spread and the dividend yield are 25.0% and 16.6%, respectively.

Those low hit ratios are not due to a change in sign of the predictor value–Panel B shows that the proportion of months in 2008 when the value of the default spread and the dividend yield remains above average equals 100% (i.e., $z_j z_{j,t} > 0$). Therefore, the poor hit ratios are driven by a change in sign in the estimated slope coefficients of the selected funds (from positive to negative). It seems that the events in 2008 have caused important structural breaks in the predictive relations that, for some reasons, did not affect the VIX. The unique nature of 2008 is confirmed by our comparative analysis. Computing hit ratios of the default spread and dividend yield during similar periods of extreme positive values over the period 1997-2007, we find a much higher stability in the predictive relation (i.e., the hit ratio amount to 87.5% and 83.3%, respectively).

To further interpret the performance of the VIX strategy, we compare its selection of hedge fund investment styles with that of the unconditional portfolio. Figure 6 plots the value of the VIX at each formation date (Panel A), next to the style composition of the unconditional portfolio (Panel B), and the VIX strategy (Panel C).

Please insert Figure 6 here

First, the VIX strategy invested more in market neutral funds than the unconditional portfolio. Market neutral funds were among the best performers in 2008, as the equal-weighted style index produced a positive 1% annual return. Second, the VIX strategy reduced its exposure to convertible funds, which fell heavily in 2008 (-21%). Thus, part of the superior performance of the VIX strategy is due to overweighting good performers (market neutral) and underweighting bad performers (convertible bonds) relative to the unconditional strategy.

During the 2008 crisis, many investors tried to withdraw their money out of the hedge fund industry. To prevent these massive outflows, hedge funds reacted by lengthening redemption notice periods, erecting gates, or creating side-pockets. To measure the costs of such liquidity constraints, we measure the performance of the different strategies between January 1997 and December 2008, after exceptionally allowing the investor to rebalance his portfolio monthly during 2008 (using all information available at the start of the month). The results in Panel C show that monthly rebalancing leads to a large improvement in performance compared to annual rebalancing (Panel A). For instance, the increase in annual alpha (Δ) is equal to 1.5%, 2.4%, and 1.7% for the unconditional, aggregate flows, and combination strategies, respectively. Overall, the liquidity constraints carried considerable costs to hedge fund investors during the crisis.

D Sensitivity Tests

So far our results indicate that the combination strategy generates higher risk-adjusted performance than all other strategies. To determine whether this conclusion is robust to alternative specifications, we perform a range of sensitivity checks reported in Table VII. First, our conclusions are robust to reducing the maximum number of funds from 75 to 50, as well as to removing this upper bound (i.e., holding the top decile portfolio even when it contains many funds). Second, repeating the analysis to include small funds (rather than imposing the AuM cutoff) leaves the relative performance of the combination strategy nearly unchanged.

In our baseline specification, we assume that when a fund stops reporting returns after being selected into the portfolio, its capital allocation is invested at the riskless rate. Therefore, we avoid look-ahead bias since we do not condition on a fund being alive during the entire year. However, funds generally disappear from the database because they are liquidated. To address this issue, we penalize any missing monthly observation with a -25% return, after which the remaining funds are invested in the riskfree asset. While the annualized alpha and mean of all strategies decrease by around 2%, the relative performance of the combination strategy slightly improves.

Most hedge fund databases receive hedge fund returns with a delay as most funds report NAV and return performance several weeks after month-end. Although large hedge fund investors may obtain return and NAV estimates from a subset of hedge funds directly, the vast majority of investors rely on commercially available data bases like ours. Examining the economic value of predictability when reporting lags are explicitly taken into account, we find that the combination strategy still significantly outperforms.

One important constraint is that of redemption notice periods–an investor who wishes to rebalance his hedge fund portfolio in December may have to give notice to the fund, typically three months in advance. To address this issue, we carry out a robustness test in which the investor has to decide, based on the information observed at the end of September, which funds to hold at December end. We find that the superior performance of the combination is robust to this change.

Finally, while our baseline specification focuses on total return predictability, funds could also be ranked based on the t-statistic of the conditional alpha in Equation (3). We document that the performance of the combination strategy based on alpha predictability is even better. This is consistent with our previous discussion that most in-sample return predictability is driven by alpha predictability. However, focusing on alpha predictability is more sensitive to the potential bias caused by omitted benchmark assets, as we use the same model to form the portfolio and evaluate subsequent performance (e.g., Carhart (1997)).

Please insert Table VII

To guard against the possibility of omitted benchmarks, in unreported results, we examine whether the alpha and Information Ratios computed using the Fung-Hsieh model change under alternative specifications. We consider the four-factor model (market, size, book-to-market, and momentum factors) and extended versions of the Fung-Hsieh model that include the Pastor and Stambaugh (2003) liquidity factor, the emerging market portfolio, and an additional equity straddle. These results remain qualitatively unchanged.

V Conclusion

The recent financial crisis has not spared the hedge fund industry. To the contrary, it has highlighted the need to search for a suitable forecasting models for hedge fund performance, as traditional measures building on past stellar performance conclusively failed. This paper embraces the challenge by developing and applying a unified framework in which to assess both in-sample and out-of-sample predictability on a fund-by-fund basis.

Using a comprehensive sample of hedge funds during 1994-2008 along with predictive variables that proxy for changing business conditions, we identify the fraction of funds in the population having future returns that are (i) negatively related; (ii) unrelated; or (iii) positively related to each of these macro variables. We find ample evidence of in-sample predictability, both across predictors and investment styles. In addition, the bulk of return predictability is attributable to predictable time-varying alphas, whereas predictable benchmark returns play only little role.

To examine whether hedge fund return predictability is evident out-of-sample, we carry out a range of tests that carefully incorporate several real-world investment constraints faced by institutional investors. We show that in a multi-fund environment, the impact of estimation risk and model uncertainty on performance can be dramatic because there is a trade-off between unconditional and predictable performance. When using poor predictive information, a hypothetical investor not only selects funds with low predictability, but may also exclude funds with high unconditional performance. We find that a conditional strategy that combines forecasts across predictors circumvents all these challenges and ultimately delivers superior performance.

The overall evidence documented here makes several contributions to the hedge fund literature. First, we provide one of the most detailed analyses to date of the statistical and economic drivers of conditional strategies' performance and highlight innovative links between in- and out-of-sample predictability. In particular, we scrutinize periods when predictor values strongly depart from their long run means and use one such period–the 2008 crisis–as a natural out-of-sample test to test the robustness of our findings. Second, while previous studies typically focus on unconditional measures of hedge fund indices, we shed light on predictability on a fund-by-fund basis. Next, our results have implications for the forecasting literature. In particular, we gain important insights on the economic value of exploiting macro conditions in investment decisions by studying the sources of the superior performance of combination strategies.

References

[1] Agarwal V., and N. Y. Naik, 2004, Risks and Portfolio Decisions Involving Hedge Funds, Review of Financial Studies 17, 63-98.

[2] Amenc N., El Bied S., and L. Martellini, 2003, Predictability in Hedge Fund Returns, Financial Analysts Journal 59, 32-46.

[3] Amihud Y., and C. M. Hurvich, 2004, Predictive Regressions: A Reduced-Bias Estimation Method, Journal of Financial and Quantitative Analysis, 813-841.

[4] Amihud Y., Hurvich C. M., and Y. Wang, 2008, Multiple-Predictor Regressions: Hypothesis Testing, Review of Financial Studies 22, 2008, 414-434.

[5] Ang A., Gorovyy S., and G. van Inwegen, 2010, Hedge Fund Leverage, Working Paper.

[6] Avramov D., 2002, Stock Return Predictability and Model Uncertainty, Journal of Financial Economics 64, 423 – 458.

[7] Avramov D., and R. Wermers, 2006, Investing in Mutual Funds When Return are Predictable, Journal of Financial Economics 81, 339-377.

[8] Avramov D., Kosowski R., Naik N. Y., and M. Teo, 2010, Hedge Funds, Managerial Skill, and Macroeconomic Variables, Journal of Financial Economics, Forthcoming.

[9] Baquero G., and M. Verbeek, 2009, A Portrait of Hedge Fund Investors: Flows, Performance and Smart Money, Working Paper.

[10] Bates J., and C. Granger, 1969, The Combination of Forecasts, Operational Research Quarterly 20, 451-468

[11] Barras L., Scaillet O., and R. Wermers, 2010, False Discoveries in Mutual Fund Performance: Measuring Luck in Estimated Alphas, Journal of Finance 65, 179-216.

[12] Bollen N. P., and V. K. Pool, 2009, Do Hedge Fund Managers Misreport Returns? Evidence from the Pooled Distribution. Journal of Finance 64, 2257-2288.

[13] Bollen N. P., and R. E. Whaley, 2009, Hedge Fund Risk Dynamics: Implications for Performance Appraisal, Journal of Finance 64, 987-1037.

[14] Carhart M. M., 1997, On Persistence in Mutual Fund Performance, Journal of Finance 52, 57-82.

[15] Choi D., Getmansky M., Henderson B., and H. Tookes, 2010, Convertible Bond Arbitrageurs as Suppliers of Capital, Review of Financial Studies 23, 2492-2522.

[16] Christopherson J. A., Ferson W., and D. A. Glassman, 1998, Conditioning Manager Alphas on Economic Information: Another Look at the Persistence of Performance, Review of Financial Studies 11, 111-142.

[17] Cochrane J. H., 2008, The Dog that Did not Bark: A Defense of Return Predictability, Review of Financial Studies 21, 1533-1575.

[18] Ding B., Getmansky M., Liang B., R. Wermers, 2009, Share Restrictions and Investor Flows in the Hedge Fund Industry, Working Paper.

[19] Elton E. J., Gruber M. J., and C. R. Blake, 1996, The Persistence of Risk-Adjusted Mutual Fund Performance, Journal of Business 69, 133-157.

[20] Fama E. F., and K. R. French, 1989, Business Conditions and Expected Returns on Stocks and Bonds, Journal of Financial Economics 25, 23-49.

[21] Ferson W., and C. Harvey, 1991, The Variation of Economic Risk Premiums, Journal of Political Economy 99, 285-315.

[22] Ferson W., and C. Harvey, 1999, Conditioning Variables and the Cross-Section of Stock Returns, Journal of Finance 54, 1325-1360.

[23] Ferson W., and M. Qian, 2004, Conditional Performance Evaluation Revisited, in Research Foundation Monograph of the CFA Institute.

[24] Fung W., and D. A. Hsieh, 2004, Hedge Fund Benchmarks: A Risk-Based Approach, Financial Analysts Journal 60, 65-80.

[25] Fung W., Hsieh D. A., Naik N., and T. Ramadorai, 2008, Hedge Funds: Performance, Risk and Capital Formation, Journal of Finance 63, 1777-1803.

[26] Ganshaw T., 2010, Hedge Funds, Humbled: The 7 Mistakes That Brought Hedge Funds to Their Knees and How They Will Rise Again. McGraw-Hill.

[27] Getmansky M., Lo A., and I. Makarov, 2004, An Econometric Model of Serial Correlation and Illiquidity in Hedge Fund Returns, Journal of Financial Economics 74, 529-609.

[28] Goyal A., and I. Welch, 2008, A Comprehensive Look at the Empirical Performance of Equity Premium Prediction, Review of Financial Studies 21, 1455-1508.

[29] Greene W. H., 2000, Econometric Analysis. 4th Edition. Prentice-Hall.

[30] Hamilton J. D., 1994, Time Series Analysis. Princeton University Press.

[31] Hendry D. F., and M. P. Clements, 2004, Pooling of Forecasts, Econometrics Journal 7, 1-31.

[32] Ilmanen A., 1995, Time-Varying Expected Returns in International Bond Markets, Journal of Finance 50, 481-506.

[33] Jagannathan R., and T. Ma, 2003, Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps, Journal of Finance 58, 1651-1683.

[34] Jagannathan R., Malakhov A., and D. Novikov, 2010, Do Hot Hands Exist Among Hedge Fund Managers? An Empirical Evaluation, Journal of Finance 65, 217-255.

[35] Jylhä P., and M. Suominen, 2010, Speculative Capital and Currency Carry Trades, Journal of Financial Economics, Forthcoming.

[36] Keim D. B., and R. F. Stambaugh, 1986, Predicting Returns in the Stock and Bond Markets, Journal of Financial Economics 17, 357-390.

[37] Kosowski R., Naik N., and M. Teo, 2007, Do Hedge Funds Deliver Alphas? A Bayesian and Bootstrap Analysis, Journal of Financial Economics, 229-264.

[38] Kosowski R., Timmermann A., Wermers R., and H. White, Can Mutual Fund "Stars" Really Pick Stocks? New Evidence from a Bootstrap Analysis, Journal of Finance 61, 2551–2596.

[39] Lhabitant F. S., 2006, Handbook of Hedge Funds. John Wiley & Sons, London.

[40] Lo A., 2002, The Statistics of the Sharpe Ratio, Financial Analysts Journal 58, 36-51.

[41] Mitchell M., and T. Pulvino, 2001, Characteristics of Risk and Return in Risk Arbitrage, Journal of Finance 61, 2135-2175.

[42] Naik N., Ramadorai T., and M. Strömqvist, 2007, Capacity Constraints and Hedge Fund Strategy Returns, European Financial Management 13, 239-256.

[43] Nelson C. R., and M. J. Kim, 1993, Predictable Stock Returns: The Role of Small Sample Bias, Journal of Finance 48, 641-661.

[44] Newey W. K., and K. D. West, 1987, A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, Econometrica 55, 703–708.

[45] Nicholls D. F., and A. L. Pope, 1988, Bias in the Estimation of Multivariate Autoregressions, Australian Journal of Statistics 30A, 296-309.

[46] Patton A., and T. Ramadorai, 2010, On the Dynamics of Hedge Fund Risk Exposures, Working Paper.

[47] Pástor L., and R. F. Stambaugh, 2003, Liquidity Risk and Expected Stock Returns, Journal of Political Economy 111, 642-8.

[48] Pástor L., and R. F. Stambaugh, 2009, Predictive Systems: Living with Imperfect Predictors, Journal of Finance 64, 1583-1628.

[49] Pesaran M. H., and A. Timmermann, 1995, Predictability of Stock Returns: Robustness and Economic Significance, Journal of Finance 50, 1201-1228.

[50] Rapach D. E., Strauss J. K., and G. Zhou, 2010, Out-of-sample Equity Premium Prediction: Combination Forecasts and Links to the Real Economy, Review of Financial Studies 23, 821-862.

[51] Stambaugh R. F., 1999, Predictive Regressions, Journal of Financial Economics 54, 375-421.

[52] Taylor S. J., Yadav P. K., and Y. Zhang, 2010, The Information Content of Implied Volatilities and Model-free Volatility Expectations: Evidence from Options Written on Individual Stocks, Journal of Banking and Finance 34, 871-88.

[53] Timmermann A., 2006, Forecast Combinations, in Elsevier Science (North-Holland): Handbook of Economic Forecasting vol. 1, ch. 4.

[54] Timmermann A., 2008, Elusive Return Predictability, International Journal of Forecasting, 1-18.

VI Appendix

A Estimation Procedure

A.1 Estimating the Slope Coefficients

Return Predictability: The Approach of Amihud, Hurvich, and Wang

For each hedge fund i in the population ($i = 1, ..., M$), we use the following predictive system:

$$\begin{aligned} r_{i,t+1} &= b_{i,0} + b_i' Z_t + u_{i,t+1}, \\ Z_{t+1} &= \theta + \Phi Z_t + \upsilon_{t+1}, \qquad t = 2, ..., T+1, \end{aligned} \tag{A1}$$

where $r_{i,t+1}$ the fund excess return (over the riskfree rate) between t and $t+1$, Z_t is the J-vector of predictors observed at time t, $b_{i,0}$ is the intercept, $b_i = [b_{i,1}, ..., b_{i,J}]'$ is the J-vector of slope coefficients, and Φ is the $J \times J$ companion matrix of the VAR(1). We denote by $u_{i,t+1}$ the fund innovation term and by υ_{t+1} the J-vector of predictor innovations between time t and $t+1$. We assume that $u_{i,t+1} = E\left(u_{i,t+1} \left| \upsilon_{t+1} \right.\right) + e_{i,t+1} = \phi_i' \upsilon_{t+1} + e_{i,t+1}$, where ϕ_i is the J-vector of innovation coefficients, and $e_{i,t+1}$ is the fund residual term (orthogonal to Z_t and υ_{t+1}). While υ_{t+1} is modelled as an i.i.d. process, we allow $e_{i,t+1}$ to be autocorrelated to account for potential illiquidity, as explained in Appendix B.2.

To construct a proxy, υ_{t+1}^c, for the unobserved J-vector υ_{t+1}, we follow the procedure initially described by Amihud and Hurvich (2004) and further developed by Amihud, Hurvich, and Wang (2008; AHW hereafter). We describe the main steps of the estimation procedure, and refer to them for further detail. First, we compute the VAR(1) estimates $\widehat{\theta}$ and $\widehat{\Phi}$. Based on these estimates, we obtain the time-series of estimated innovation vector, $\widehat{\upsilon}_{t+1}$, from which we compute the $J \times J$ innovation covariance matrix, denoted by $\widehat{\Sigma}_\upsilon$. Second, we use $\widehat{\theta}$, $\widehat{\Phi}$, and $\widehat{\Sigma}_\upsilon$ to correct for the small-sample bias in $\widehat{\Phi}$ using the analytical formula proposed by Nicholls and Pope (1988):

$$\widehat{bias}(\widehat{\Phi}) = -\frac{1}{T}\widehat{\Sigma}_\upsilon \left[\left(I_J - \widehat{\Phi}'\right)^{-1} + \widehat{\Phi}'\left(I_J - \widehat{\Phi}'^2\right)^{-1} + \sum_{j=1}^{J} \lambda_j \left(I_J - \lambda_j \widehat{\Phi}'\right)^{-1} \right] \widehat{\Sigma}_Z^{-1}, \tag{A2}$$

where T is the number of observations, I_J is a $J \times J$ identity matrix, and λ_j denotes the j^{th} eigenvalue of $\widehat{\Phi}'$, and $\widehat{\Sigma}_Z$ is the $J \times J$ covariance matrix of Z_t computed using the following formula: $vec(\Sigma_\upsilon) = (I_{J^2} - A)vec(\Sigma_Z)$, where vec is the vec operator, I_{J^2}

is a $J^2 \times J^2$ identity matrix, and $A = (\Phi \otimes \Phi)$ (see Hamilton (1994), p. 265). The bias formula , $\widehat{bias}(\widehat{\Phi})$, is estimated iteratively. In each iteration k ($k = 2, ..., K$), we use the following updating scheme: $\widehat{\Phi}_{(k)} = \widehat{\Phi}_{(k-1)} - \widehat{bias}(\widehat{\Phi})_{(k)}$ and $\widehat{\theta}_{(k)} = \left(I - \widehat{\Phi}_{(k)}\right)\overline{Z}$, where $\overline{Z}$ denotes the sample mean, and $\widehat{\Sigma}_{\upsilon(k)}$ is obtained using the updated estimates, $\widehat{\Phi}_{(k)}$ and $\widehat{\theta}_{(k)}$. As in AHW, the number of iterations, K, is set equal to 10. Third, we use the final bias-corrected VAR(1) estimates, $\widehat{\theta}^c$ and $\widehat{\Phi}^c$, to construct the proxy, υ_t^c:

$$\upsilon_t^c = Z_{t+1} - \widehat{\theta}^c + \widehat{\Phi}^c Z_t, \qquad t = 2, ..., T+1. \tag{A3}$$

To compute the J-vector of bias-corrected estimated slope coefficients, $\widehat{b}_i$, we replace υ_t with υ_t^c in Equation (11), and regress the fund return, $r_{i,t+1}$, on an augmented vector x_t including $2{\cdot}J + 1$ explanatory variables: $x_t = [1, Z_t', \upsilon_{t+1}^{c\prime}]'$. Replacing $u_{i,t+1}$ with $\phi_i' \upsilon_{t+1} + e_{i,t+1}$ in Equation (A1) and υ_{t+1} with $\upsilon_{t+1}^c + (\widehat{\theta}^c - \theta) + \left(\widehat{\Phi}^c - \Phi\right) Z_t$, we can determine the remaining bias in the J-vector $\widehat{b}_i$: $bias(\widehat{b}_i) = E\left(\widehat{\Phi}^c - \Phi\right)' \phi_i$. We show that $\widehat{b}_i$ is nearly unbiased since the estimated companion matrix, $\widehat{\Phi}^c$, is corrected for the small sample bias and gets very close to the true value, Φ.[19]

From Return to Alpha Predictability

AHW focus on time series predictive regressions. Here, we extend their approach to control for small-sample bias in the asset pricing regression with time varying alpha formulated in equation (3):

$$r_{i,t+1} = a_{i,0} + a_i' Z_t + \beta_i' f_{t+1} + \epsilon_{i,t+1}, \tag{A4}$$

where $a_{i,0}$ is the intercept, $a_i = [a_{i,1}, ..., a_{i,J}]'$ is the J-vector of alpha slope coefficients, and β_i the K-vector of fund exposure to the K risk factors, f_{t+1}, and $\epsilon_{i,t+1}$ is the new innovation term (orthogonal to Z_t and F_{t+1}). Projecting the J-vector of predictor innovation, υ_{t+1}, onto the space spanned by f_{t+1}, we obtain a new innovation vector denoted by ω_{t+1}, i.e., $\omega_{t+1} = \upsilon_{t+1} - q_\upsilon - Q_\upsilon f_{t+1}$, where q_υ is a J-vector and Q_υ a $J \times K$ coefficient matrix. As in Equation (A1), the new innovation terms are correlated: $\epsilon_{i,t+1} = E\left(\epsilon_{i,t+1} \,|\omega_{t+1}\right) + \varrho_{i,t+1} = \psi_i' \omega_{t+1} + \varrho_{i,t+1}$, where ψ_i is the J-vector of innovation coefficients, and $\varrho_{i,t+1}$ is the fund residual term (orthogonal to Z_t, f_{t+1}, and ω_{t+1}).

Similar to Equation (11), the small-sample bias in the J-vector of OLS-estimated alpha slope coefficients, $\widehat{a}_i^{ols}$, can be eliminated if we include the J-vector ω_{t+1} as an

[19] As shown by Nicholls and Pope (1988), the estimated bias, $\widehat{bias}(\widehat{\Phi})$, shown in Equation (A2) is very close from the true (but unobservable) bias, $bias(\widehat{\Phi})$, since $bias(\widehat{\Phi}) = \widehat{bias}(\widehat{\Phi}) + O(T^{-\frac{3}{2}})$.

additional vector of explanatory variables, i.e.,

$$r_{i,t+1} = a_{i,0} + a_i' Z_t + \beta_i' f_{t+1} + \psi_i' \omega_{t+1} + \varrho_{i,t+1}, \tag{A5}$$

because the orthogonality condition holds, i.e., $E(\varrho_i | X) = 0$, where $\varrho_i = [\varrho_{i,1}, ..., \varrho_{i,T+1}]'$, $X = [x_1, ..., x_{T+1}]'$, and $x_t = [Z_{t-1}', \upsilon_t', f_t']'$. Of course, we cannot observe the true innovation vector thus we have to find a proxy for it denoted by υ_{t+1}^c. To proxy for the unobservable vector, ω_{t+1}, we take the J-vector of innovation, υ_{t+1}^c, computed in Equation (A3) and regress it on the factor returns, f_{t+1}, and a constant to obtain

$$\omega_{t+1}^c = \upsilon_{t+1}^c - \widehat{q}_\upsilon - \widehat{Q}_\upsilon f_{t+1}. \tag{A6}$$

To compute the J-vector of bias-corrected estimated alpha slope coefficients, $\widehat{a}_i$, we replace ω_{t+1} with ω_{t+1}^c in Equation (A5), and regress the fund return, $r_{i,t+1}$, on an augmented vector x_t including $2J+K+1$ explanatory variables variables: $x_t = [1, Z_t', f_{t+1}', \omega_{t+1}^{c\prime}]'$. We can replace $\epsilon_{i,t+1}$ with $\psi_i' \omega_{t+1} + \varrho_{i,t+1}$ in Equation (A4), where

$$\begin{aligned} \omega_{t+1} &= \upsilon_{t+1} - q_{\upsilon,0} - Q_\upsilon f_{t+1} \\ &= \omega_{t+1}^c + (\widehat{\theta}^c - \theta_j) + \left(\widehat{\Phi}^c - \Phi\right) Z_t + (\widehat{q}_\upsilon - q_\upsilon) + \left(\widehat{Q}_\upsilon - Q_\upsilon\right)' f_{t+1}, \end{aligned} \tag{A7}$$

to get an expression for the remaining bias that is similar to the one obtained for the return predictive regression, i.e., $bias(\widehat{a}_i) = E\left(\widehat{\Phi}^c - \Phi\right)' \psi_i$.

From Monthly to Quarterly Return Predictability

Here, we extend the AHW approach to assess hedge fund predictability at the quarterly horizon. Let us denote by $r_{i,t:t+1}$, $r_{i,t+1:t+2}$, and $r_{i,t+2:t+3}$ the (excess) return of fund i between time t and $t+1$, $t+1$ and $t+2$, and $t+2$ and $t+3$, respectively. Based on Equation (A1), we can write:

$$\begin{aligned} r_{i,t:t+1} &= b_{i,0} + b_i' Z_t + \phi_i' \upsilon_{t+1} + e_{i,t+1}, \\ r_{i,t+1:t+2} &= b_{i,0} + b_i' Z_{t+1} + \phi_i' \upsilon_{t+2} + e_{i,t+2} \\ &= b_{i,0} + b_i' (\theta + \Phi Z_t + \upsilon_{t+1}) + \phi_i' \upsilon_{t+2} + e_{i,t+2}, \\ r_{i,t+2:t+3} &= b_{i,0} + b_i' Z_{t+2} + \phi_i' \upsilon_{t+3} + e_{i,t+3}, \\ & \quad b_{i,0} + b_i' (\theta + \Phi(\theta + \Phi Z_t + \upsilon_{t+1}) + \upsilon_{t+2}) + \phi_i' \upsilon_{t+3} + e_{i,t+3}. \end{aligned} \tag{A8}$$

Ignoring compounding effects, the monthly growth rate of fund i over the next quarter, denoted by $r^q_{i,t+3}$, is equal to $\frac{1}{3}(r_{i,t:t+1}+r_{i,t+1:t+2}+r_{i,t+2:t+3})$. After grouping the different terms in Equation (A8), we can write $r^q_{i,t+3}$ as

$$r^q_{i,t+3} = b^q_{i,0} + b^{q\prime}_i Z_t + \phi^{q\prime}_{i,1} \upsilon_{t+1} + \phi^{q\prime}_{i,2} \upsilon_{t+2} + \phi^{q\prime}_{i,2} \upsilon_{t+3} + e^q_{i,t+3}, \tag{A9}$$

where $b^q_{i,0} = \frac{1}{3}(3b_{i,0} + 2b_i'\theta + 2b_i'\Phi\theta)$, $b^q_i = \frac{1}{3}(b_i + b_i\Phi' + b_i\Phi^{2\prime})$, $\phi^q_{i,1} = \frac{1}{3}(\phi_i + b_i + b_i\Phi')$, $\phi^q_{i,2} = \frac{1}{3}(\phi_i + b_i)$, $\phi^q_{i,3} = \frac{1}{3}\phi_i$, and $e^q_{i,t+3} = \frac{1}{3}(e_{i,t+1} + e_{i,t+2} + e_{i,t+3})$. Equation (A9) reveals that we need to include three lags of the J-vector of innovation, υ_{t+1}, υ_{t+2}, and υ_{t+3}, to control for small sample bias at a quarterly horizon. To apply the method of Amihud, Hurvich, and Wang, we use the proxy υ^c_{t+1} computed in Equation (A3). Then, we construct $r^q_{i,t:t+3}$ using overlapping monthly observations and regress it on the $4 \cdot J + 1$-vector, x_t, of explanatory variables variables: $x_t = [1, Z_t', \upsilon^{c\prime}_{t+1}, \upsilon^{c\prime}_{t+2}, \upsilon^{c\prime}_{t+3}]$ to obtain the J-vector of bias-corrected estimated slope coefficients, $\widehat{b}^q_i$ at the quarterly horizon.

A.2 Estimating the Slope Coefficient t-statistic (and p-values)

We extend the approach of AHW to incorporate hedge fund illiquidity when computing the variance of the estimated bias-corrected slope coefficient, $\widehat{b}_{i,j}$, associated with each predictor j $(j = 1, ..., J)$. Following AHW, we can write the estimated variance of $\widehat{b}_{i,j}$ as

$$\widehat{var}(\widehat{b}_{i,j}) = \sum_{i=1}^{J} \sum_{k=1}^{J} \widehat{\phi}_{i,j} \widehat{\phi}_{k,j} \widehat{cov}(\widehat{\rho}_{ij}, \widehat{\rho}_{kj}) + \widehat{var}_{aug}(\widehat{b}_{i,j}), \tag{A10}$$

where $\widehat{\rho}_{i,j}$ denotes the i^{th} row-j^{th} column element of the estimated companion matrix, $\widehat{\Phi}^c$, and $\widehat{\phi}_{i,j}$ is the j^{th} element of the J-vector of estimated innovation coefficients, $\widehat{\phi}_i$. The terms $\widehat{cov}(\widehat{\rho}_{ij}, \widehat{\rho}_{kj})$ are read from the estimated covariance matrix of the VAR(1). Finally, $\widehat{var}_{aug}(\widehat{b}_{i,j})$ is the $(j+1)$ diagonal element of the $(2J+1) \times (2J+1)$ estimated covariance matrix of the augmented regression, $\widehat{V}_{aug}(\widehat{b}_{i,0}, \widehat{b}_i, \widehat{\phi}_i) = (X'X)^{-1} \left(X'\widehat{V}_i X\right)^{-1} (X'X)^{-1}$, where $X_{(T\times(2\cdot J+1))} = [x_1, ..., x_T]'$, $x_1 = [1, Z_1', \upsilon^{c\prime}_2]'$, and $\widehat{V}_i$ is the $T \times T$ estimated covariance matrix of the fund residual vector $e_i = [e_{i,2}, ..., e_{i,T+1}]'$.

To account for autocorrelation caused by potential illiquidity, we estimate V_i using an AR specification. Based on empirical evidence (discussed in Section II.C.2 of the paper), we use a two-month lag, i.e.,

$$e_{i,t+1} = \rho_{i,1} e_{i,t} + \rho_{i,2} e_{i,t-1} + \xi_{i,t+1}. \tag{A11}$$

After estimating the coefficients in Equation (A11) from fund i estimated innovation, $\widehat{e}_{i,t+1}$, we compute $\widehat{V}_i$ as $\widehat{var}(\widehat{\xi}_{i,t+1})\left(\widehat{\Psi}_i'\widehat{\Psi}_i\right)^{-1}$, where $\widehat{\Psi}_i$ is defined as

$$\begin{bmatrix} \left(\frac{(1+\widehat{\rho}_{i,2})\left[(1-\widehat{\rho}_{i,2}^2)-\widehat{\rho}_{i,1}^2\right]}{1-\widehat{\rho}_{i,2}}\right)^{\frac{1}{2}} & 0 & 0 & 0 & ... & 0 & 0 & 0 \\ \frac{\widehat{\rho}_{i,1}(1-\widehat{\rho}_{i,1}^2)^{\frac{1}{2}}}{1-\widehat{\rho}_{i,2}} & \left(1-\widehat{\rho}_{i,2}^2\right)^{\frac{1}{2}} & 0 & 0 & ... & 0 & 0 & 0 \\ -\widehat{\rho}_{i,2} & -\widehat{\rho}_{i,1} & 1 & & ... & 0 & 0 & 0 \\ ... & ... & ... & ... & ... & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & ... & -\widehat{\rho}_{i,2} & -\widehat{\rho}_{i,1} & 1 \end{bmatrix}, \quad \text{(A12)}$$

and $\widehat{var}(\widehat{\xi}_{i,t+1})$ is the estimated variance of the innovation term in Equation (A11) (see Greene (2000)).

A.3 The Proportions of Unpredictable and Predictable Funds

We briefly outline the methodology proposed by BSW to estimate the proportions of funds whose returns are unrelated ($b_{i,j} = 0$), negatively related ($b_{i,j} < 0$), and positively related ($b_{i,j} > 0$) to predictor j ($j = 1, ..., J$). These proportions are denoted by $\pi_R^0(j)$, $\pi_R^-(j)$, and $\pi_R^+(j)$, respectively. Importantly, the BSW approach allows to measure true predictability, as it excludes funds that exhibit predictability only by luck, i.e., funds with significant estimated slope coefficients, $\widehat{b}_{i,j}$, while the true slope coefficient, $b_{i,j}$, equals zero by luck (these funds are "false discoveries" using their terminology).

The only required input is the M-vector of p-values associated with the estimated slope coefficient, $\widehat{b}_{i,j}$ of each fund i in the population ($i = 1, ..., M$). To this end, we use the estimated variance in Equation (A10), and compute the slope t-statistic of each fund i as $t(\widehat{b}_{i,j}) = \widehat{b}_{i,j}/\widehat{var}(\widehat{b}_{i,j})^{\frac{1}{2}}$. Then, we follow AHW and compute the fund p-value as $p(\widehat{b}_{i,j}) = 2(1 - F_N(\left|t(\widehat{b}_{i,j})\right|))$, where F_N is the cumulative function of a normal distribution.

We begin with the estimation of $\pi_R^0(j)$. Since unpredictable funds satisfy the null hypothesis $H_0 : b_{i,j} = 0$, their p-values are uniformly distributed over the interval $[0, 1]$. To recover this uniform distribution, we simply take a sufficiently high threshold λ^* beyond which the vast majority of p-values come from the unpredictable funds. After measuring the proportion $\widehat{w}(j, \lambda^*)$ of p-values above λ^* and multiplying it by $1/(1-\lambda^*)$, we obtain an estimate of the proportion of unpredictable funds:

$$\widehat{\pi}_R^0(j) = \widehat{w}(j, \lambda^*) \cdot (1 - \lambda^*)^{-1}. \quad \text{(A13)}$$

To choice of optimal threshold λ^* is based on the bootstrap procedure proposed by Storey (2002) and Storey, Taylor, and Siegmund (2004). This resampling approach consists of minimizing an estimate of the Mean-Squared Error (MSE) of $\widehat{\pi}_R^0(j,\lambda)$, defined as $E(\widehat{\pi}_R^0(j,\lambda) - \pi_R^0(j))^2$. First, we compute $\widehat{\pi}_R^0(j,\lambda)$ using Equation (A13) across a range of λ values ($\lambda = 0.30, 0.35, ..., 0.70$). Second, for each possible value of λ, we form 1,000 bootstrap replications of $\widehat{\pi}_R^0(j,\lambda)$ by drawing with replacement from the $M \times 1$ vector of fund p-values. These are denoted by $\widehat{\pi}_R^{b,0}(\lambda)$, for $b = 1, ..., 1,000$. Third, we compute the estimated MSE for each possible value of λ:

$$\widehat{MSE}(j,\lambda) = \frac{1}{1,000}\sum_{b=1}^{1,000}\left[\widehat{\pi}_R^{b,0}(j,\lambda) - \min_{\lambda}\widehat{\pi}_R^0(j,\lambda)\right]^2. \tag{A14}$$

We choose λ^* such that $\lambda^* = \arg\min_\lambda \widehat{MSE}(j,\lambda)$.

To estimate the proportions of predictable funds, $\pi_R^-(j)$ and $\pi_R^+(j)$, we use a similar approach. First, we compute $\widehat{\pi}_R^-(\gamma)$ across a range of significance levels γ ($\gamma = 0.30, 0.35, ..., 0.50$): $\widehat{\pi}_R^-(j) = \widehat{S}_\gamma^-(j) - \widehat{F}_\gamma^-(j)$, where $\widehat{S}_\gamma^-(j)$ is the observed proportions of significant funds with negative estimated slope coefficient, $\widehat{b}_{i,j}$, (at the significance level γ), and $\widehat{F}_\gamma^-(j) = \widehat{\pi}_R^0(j)\cdot\gamma/2$ is the estimated proportion of "false discoveries" (see BSW). Second, we form 1,000 bootstrap replications of $\widehat{\pi}_R^-(\gamma)$ for each possible value of γ. These are denoted by $\widehat{\pi}_R^{b,-}(j,\gamma)$, for $b = 1, ..., 1,000$. Third, we compute the estimated MSE for each possible value of γ:

$$\widehat{MSE}^-(j,\gamma) = \frac{1}{1,000}\sum_{b=1}^{1,000}\left[\widehat{\pi}_R^{b,-}(j,\gamma) - \max_{\gamma}\widehat{\pi}_R^-(j,\gamma)\right]^2. \tag{A15}$$

We choose γ^- such that $\gamma^- = \arg\min_\gamma \widehat{MSE}^-(j,\gamma)$. We use the same data-driven procedure for $\widehat{\pi}_R^+(j,\gamma)$ to determine $\gamma^+ = \arg\min_\gamma \widehat{MSE}^+(j,\gamma)$. If $\min_\gamma \widehat{MSE}^-(j,\gamma) < \min_\gamma \widehat{MSE}^+(j,\gamma)$, we set $\widehat{\pi}_R^-(j) = \widehat{\pi}_R^-(j,\gamma^-)$. To preserve the equality $1 = \pi_R^0 + \pi_R^- + \pi_R^+$, we set $\widehat{\pi}_R^+(j) = 1 - \widehat{\pi}_R^0(j) - \widehat{\pi}_R^-(j)$. Otherwise, we set $\widehat{\pi}_R^+(j) = \widehat{\pi}_{r,A}^+(j,\gamma^+)$ and $\widehat{\pi}_R^-(j) = 1 - \widehat{\pi}_R^0(j) - \widehat{\pi}_R^+(j)$.

We also compute the estimated proportion of predictable funds, $\widehat{\pi}_R^{Joint}$, when all predictors are considered simultaneously (shown in the final column of Table II). First, we compute the Wald test suggested by AHW: $w(\widehat{b}_i) = \widehat{b}_i'\widehat{V}\left(\widehat{b}_i\right)^{-1}\widehat{b}_i$, where $\widehat{V}\left(\widehat{b}_i\right)$ is the $J \times J$ estimated covariance matrix of the J-vector $\widehat{b}_i$. While the diagonal elements of $\widehat{V}\left(\widehat{b}_i\right)$ are given by Equation (A10), AHW show that a similar expression holds for the covariance terms, $\widehat{cov}(\widehat{b}_{i,j}, \widehat{b}_{i,k})$. Then, we compute the p-value associated with this joint test as $p(\widehat{b}_i) = 1 - F_N(w(\widehat{b}_i))$, where F_N is the cumulative function of a χ^2 distribution with J degrees of freedom. Finally, we use Equation (A13) to obtain our estimate.

A.4 Estimating the Conditional t-statistic

As discussed in Section B.1, the conditional strategy selects funds with the highest conditional t-statistic, $t(\widehat{\mu}_{i,t}) = \widehat{\mu}_{i,t}/\widehat{var}(\widehat{\mu}_{i,t})^{\frac{1}{2}}$, where $\widehat{\mu}_{i,t}$ is the fund estimated conditional mean and $\widehat{var}(\widehat{\mu}_{i,t})$ denotes its estimated variance (see Equation (6)). In this section, we explain how to compute the conditional t-statistic both in the single- and multi-predictor case.

In the case of the single-predictor strategy that uses predictor j $(j+1, ..., J)$, we can write $\widehat{\mu}_{i,t}(j)$ and $\widehat{var}(\widehat{\mu}_{i,t}(j))^{\frac{1}{2}}$ as

$$\widehat{\mu}_{i,t}(j) = \widehat{b}_{i,0} + \widehat{b}_{i,j} Z_{j,t}, \qquad \widehat{var}(\widehat{\mu}_{i,t}(j)) = X_t' \widehat{V}\left(\widehat{b}_{i,0}, \widehat{b}_{i,j}\right) X_t, \tag{A16}$$

where $\widehat{b}_{i,0}$ and $\widehat{b}_{i,j}$ denote the bias-corrected estimated intercept and slope coefficient, $Z_{j,t}$ is the predictor value at time t, $X_t = [1, Z_{j,t}]'$, and $\widehat{V}\left(\widehat{b}_{i,0}, \widehat{b}_{i,j}\right)$ is the 2×2 estimated covariance matrix of the regression coefficients. To estimate $\widehat{b}_{i,0}$ and $\widehat{b}_{i,j}$, we use the simplified approach proposed by Amihud and Hurvich (2004, AH hereafter) for single-predictor regressions. First, we estimate the AR(1) model for predictor j : $Z_{j,t+1} = \widehat{\theta}_j + \widehat{\rho}_j Z_{j,t} + \widehat{\upsilon}_{j,t+1}$. Second, we replace the unobservable innovation term, $\upsilon_{j,t+1}$, with $\upsilon^c_{j,t+1} = Z_{j,t+1} - \widehat{\theta}^c_j - \widehat{\rho}^c_j Z_{j,t}$, where $\widehat{\rho}^c_j$ is the second-order bias corrected autocorrelation coefficient suggested by AH. Specifically,

$$\begin{aligned} \widehat{\rho}^c_j &= \widehat{\rho}_j + (1 + 3\widehat{\rho}_j)/T + 3(1 + 3\widehat{\rho}_j)/T^2, \\ \widehat{\theta}^c_j &= 1/(1 - \widehat{\rho}_j)\overline{Z}_j, \end{aligned} \tag{A17}$$

where T denotes the number of return observations, and $\overline{Z}_j$ denotes the sample mean. Third, we regress the hedge fund return, $r_{i,t+1}$, on the augmented vector x_t including 3 explanatory variables, $x_t = [1, Z_{j,t}, \upsilon^c_{j,t+1}]'$.

To compute $\widehat{V}\left(\widehat{b}_{i,0}, \widehat{b}_{i,j}\right)$ we follow AH and write the estimated variances of $\widehat{b}_{i,0}$ and $\widehat{b}_{i,j}$, as well as their estimated covariance as

$$\begin{aligned} \widehat{var}(\widehat{b}_{i,0}) &= \widehat{\phi}^2_{i,j} \widehat{var}(\widehat{\theta}^c_j) + \widehat{var}_{aug}(\widehat{b}_{i,0}), \\ \widehat{var}(\widehat{b}_{i,j}) &= \widehat{\phi}^2_{i,j} \widehat{var}(\widehat{\rho}^c_j) + \widehat{var}_{aug}(\widehat{b}_{i,j}) \\ \widehat{cov}(\widehat{b}_{i,0}, \widehat{b}_{i,j}) &= \widehat{\phi}^2_{i,j} \widehat{cov}(\widehat{\theta}^c_j, \widehat{\rho}^c_j) + \widehat{cov}_{aug}(\widehat{b}_{i,0}, \widehat{b}_{i,j}), \end{aligned} \tag{A18}$$

where $\widehat{var}(\widehat{\theta}_j^c)$, $\widehat{var}(\widehat{\rho}_j^c)$, and $\widehat{cov}(\widehat{\theta}_j^c, \widehat{\rho}_j^c)$ are read from the estimated covariance matrix of the AR(1) model. The terms $\widehat{var}_{aug}(\widehat{b}_{i,0})$ and $\widehat{var}_{aug}(\widehat{b}_{i,j})$ are the first two diagonal elements of the 3×3 estimated covariance matrix of the coefficients of the augmented regression, $\widehat{V}_{aug}(\widehat{b}_{i,0}, \widehat{b}_{i,j}, \widehat{\phi}_{i,j})$, while $\widehat{cov}_{aug}(\widehat{b}_{i,0}, \widehat{b}_{i,j})$ is the first row-second column off-diagonal term. Note that $\widehat{var}_{aug}(\widehat{b}_{i,j})$ is the single-predictor counterpart to the estimated variance shown in Equation (A10).

In the multi-predictor case ($J > 1$), we follow the approach outlined in Appendix B.2 to compute the J-vector of estimated slope coefficients, $\widehat{b}_i$, and their covariance matrix, $\widehat{V}\left(\widehat{b}_i\right)$. The estimates for $\widehat{var}(\widehat{b}_{i,0})$ and $\widehat{cov}(\widehat{b}_{i,0}, \widehat{b}_i)$ are slightly more complicated than those shown in Equation (A18) because the estimated intercept, $\widehat{b}_{i,0}$, is a function of the J-vector of intercept, $\widehat{\theta}^c$, i.e., $\widehat{b}_{i,0} = b_{i,0} + \phi_i'(\widehat{\theta}^c - \theta)$. However, the logic behind the approach remains unchanged.

B Description of Hedge Fund Return Data

We evaluate hedge fund performance using monthly net-of-fee returns of live and dead hedge funds using a new data base that aggregates data reported in the BarclayHedge, TASS, HFR, CISDM and MSCI. The union of databases represents the largest known dataset of hedge funds to date. A key advantage of this database is that it allows to reduce selection bias. Since not all hedge funds report to commercial hedge fund data bases, our aggregate data base is likely to be closer to the true unobserved population of hedge funds. Our initial fund universe contains 33,244 live and dead hedge funds.

To create this data set, we control for a number of potential biases. First, we are careful to remove duplicate funds that exist in several databases, as well as funds with multiple shareclasses.[20] Second, prior to 1994, databases are subject to survivorship bias, since they did not keep track of the hedge funds that died. This bias is mitigated by only examining the period from January 1994 onwards. Third, funds often report return data prior to their listing date in the database, thereby creating a backfill bias. Since well-performing funds have strong incentives to list, the backfilled returns are usually higher than the non-backfilled returns. To address this issue, we exclude the first 12 months of data for each fund.

To allow a detailed interpretation of predictability results per strategy we choose to group funds into 10 categories, finer than examined in past work: long-short equity, equity market neutral, managed futures, global macro, emerging markets, convertible

[20] We use two main approaches to identify duplicates between the databases. The first is based on a string comparison of fund names. The second compares fund returns based on their mean difference and correlations. When multiple shareclasses are identified we use the oldest US dollar shareclass.

arbitrage, event driven, fixed income, fund of funds, and multi-strategy.

Long-short equity funds take long and short positions in undervalued and overvalued stocks, respectively, and reduce systematic risks in the process. Equity market neutral funds are similar to long-short equity funds in that they take long and short positions, but differ in that they typically follow more high frequency signals and systematic trading strategies.

Managed futures funds include Commodity Trading Advisors (CTAs) and share some similarities with macro funds in that they use relatively liquid instruments such as futures and often pursue directional strategies. However, they differ in the type of signals used which tend to be more high frequency and quantitative in nature. Macro funds use a discretionary investment strategy that depends on global macroeconomic variables and often reflects a medium- to long-term outlook. Emerging markets funds pursue a range of (historically mostly long-only) strategies in emerging markets.

Convertible arbitrage funds exploit mispricing in the convertible bond market such as underpriced implied volatility, for example. Event-driven funds, which include merger arbitrage funds, monitor corporate events and restructurings and employ multiple strategies usually involving investments in opportunities created by significant transactional events (such as spin-offs, mergers and acquisitions, bankruptcy reorganizations, recapitalizations, and share buybacks). Fixed income (relative value) funds follow a range of spread strategies in different parts of the fixed income market to benefit from relative mispricing related to credit risk or the shape of the yield curve.

Multi-strategy funds are similar to funds of funds in that they attempt to achieve diversification across strategies. But unlike the latter, they often follow more specialized strategies that, given fewer constraints on redemptions and inflows than fund of funds, can respond quickly to tactical signals.The final investment category is funds of funds.

As Table I of the paper shows, the majority of funds belong to the funds of funds (2,034) and long-short equity (1,393) categories, followed by managed futures (830), multi-strategy (687), emerging markets (389), event-driven (277), fixed income (256), macro (237), equity market neutral (217), and convertible arbitrage (201).

Table I

Descriptive Statistics

Panel A shows, for each investment category, the number of funds and their relative importance in the population (in parentheses), as well as the fund cross-sectional median and the 25-75% quantiles (in brackets) of the annualized excess mean (over the riskfree rate) and standard deviation, the skewness, and the kurtosis. Panel B displays the monthly mean and standard deviation, as well as the first-order autocorrelation and correlation matrix of the four predictive variables used to forecast hedge fund returns. Panel C displays the annualized excess mean and standard deviation, as well as the correlation matrix of the Fung and Hsieh (2004) seven risk factors. All statistics are computed using monthly observations between January 1994 and December 2007.

Panel A Fund Excess Returns

	Number(%)	Mean(Ann.)	Std(Ann.)	Skewness	Kurtosis
All Funds	7,991(100)	5.2%[1.2,9.5]	10.2%[6.4,16.6]	-.03 [-.41,.33]	3.7 [2.9,5.0]
Long-Short	1,393(17.4)	7.1 [2.7,11.5]	12.9 [8.8,17.9]	.03 [-.24,.35]	3.6 [2.8,4.5]
Market Neutral	217(2.7)	2.6 [-0.8,6.6]	7.4 [5.1,9.5]	-.07 [-.37,.29]	3.6 [2.9,5.1]
Managed Futures	830(10.4)	3.2 [-1.3,8.6]	15.4 [9.6,21.2]	.25 [.00,.54]	3.4 [2.8,4.4]
Global Macro	237(3.0)	5.5 [1.5,9.1]	13.3 [8.8,18.1]	.21 [-.20,.53]	3.7 [3.1,4.8]
Emerging Markets	389(4.9)	8.8 [3.7,19.1]	17.4 [10.2,23.1]	-.08 [-.33,.20]	3.3 [2.7,4.4]
Convertible Arb.	201(2.5)	3.3 [0.3,6.2]	5.5 [3.8,8.5]	-.23 [-.66,.13]	4.2 [3.4,5.8]
Event-Driven	277(3.5)	6.2 [2.8,10.7]	7.6 [5.1,11.7]	-.23 [-.71,.26]	4.6 [3.5,6.6]
Fixed Income	256(3.2)	3.0 [-0.3,5.7]	7.2 [3.9,9.6]	-.37 [-1.21,.12]	5.0 [3.4,8.2]
Funds of Funds	2,034(24.5)	3.9 [0.5,6.3]	7.0 [4.3,9.2]	-.31 [-.65,.06]	3.9 [3.1,5.3]
Multi-Strategy	687(8.6)	5.9 [2.2,10.6]	13.8 [8.2,22.2]	.16 [-.12,.46]	3.3 [2.6,4.6]

Panel B Predictors

				Correlation matrix		
	Mean(Mon.)	Std(Mon.)	Autocorr.	Dividend	Volatility	Agg. Flows
Default Spread	0.8%	0.2%	0.95	-.26	.30	.19
Dividend Yield	2.0	0.4	0.97		-.48	-.26
VIX (Volatility)	19.5	6.7	0.84			.01
Aggregate Flows	0.9	2.0	0.23			

Panel C Risk Factor Returns

			Correlation matrix					
	Mean(Ann.)	Std(Ann.)	Size	Term	Def.	T. Bond	T. Cur.	T. Com.
Equity Market	7.2%	13.9%	-.06	-.11	.30	-.14	-.12	-.09
Equity Size	-2.7	13.1		-.15	.20	-.05	.02	-.02
Bond Term	2.4	7.1			-.33	.06	.14	.08
Bond Default	2.2	4.1				-.12	-.15	-.12
Trend Bond	-17.2	51.6					.16	.16
Trend Currency	-3.6	64.8						.26
Trend Commodity	-8.8	46.1						

Table II

Return versus Alpha Predictability

We measure hedge fund predictability in the entire population (Panel A), and across investment categories (Panels B to K). In the first row of each panel (Return), we report, for each predictor (default spread, dividend yield, VIX, and aggregate flows), the estimated proportions of funds in the population exhibiting return predictability, $\widehat{\pi}_R^-$ and $\widehat{\pi}_R^+$. We also report the cross-fund median, $\bar{b}_j$, and 25-75% quantiles (in brackets) of the (bias-corrected) estimated slope coefficient shown in Equation (1). Each fund coefficient is standardized (by multiplying the initial estimate by the predictor standard deviation) so that it corresponds to the change in the fund monthly return for a one standard deviation increase in the predictor value. The final column (Joint) shows the estimated proportion of predictable funds using all predictors simultaneously, $\widehat{\pi}_R^{Joint}$. The inputs used to compute $\widehat{\pi}_R^{Joint}$ are the Wald tests of joint significance computed for each fund in the population. In the second row of each panel (Alpha), we repeat the same procedure using the (bias-corrected) estimated slope coefficients in Equation (3) to measure alpha predictability. All statistics are computed using monthly data between January 1994 and December 2007.

	Default Spread			Dividend Yield			VIX(Volatility)			Aggregate Flows			Joint
	Prop.(%)		Slope Coeff.	Prop.(%)		Slope Coeff.	Prop.(%)		Slope Coeff.	Prop.(%)		Slope Coeff.	Prop.(%)
Return	$\widehat{\pi}_R^-$	$\widehat{\pi}_R^+$	$\bar{b}_j[.25,.75]$	$\widehat{\pi}_R^-$	$\widehat{\pi}_R^+$	$\bar{b}_j[.25,.75]$	$\widehat{\pi}_R^-$	$\widehat{\pi}_R^+$	$\bar{b}_j[.25,.75]$	$\widehat{\pi}_R^-$	$\widehat{\pi}_R^+$	$\bar{b}_j[.25,.75]$	$\widehat{\pi}_R^{Joint}$
Alpha	$\widehat{\pi}_\alpha^-$	$\widehat{\pi}_\alpha^+$	$\bar{a}_j[.25,.75]$	$\widehat{\pi}_\alpha^-$	$\widehat{\pi}_\alpha^+$	$\bar{a}_j[.25,.75]$	$\widehat{\pi}_\alpha^-$	$\widehat{\pi}_\alpha^+$	$\bar{a}_j[.25,.75]$	$\widehat{\pi}_\alpha^-$	$\widehat{\pi}_\alpha^+$	$\bar{a}_j[.25,.75]$	$\widehat{\pi}_\alpha^{Joint}$
							Panel A All Funds						
Return	2.8	13.4	.07[-.21,.34]	22.2	4.1	-.11[-.43,.21]	22.7	3.1	-.14[-.50,.20]	33.2	0.0	-.23[-.56,.00]	60.5
Alpha	7.5	16.7	.04[-.22,.32]	21.7	3.4	-.11[-.44,.19]	21.0	4.9	-.10[-.46,.23]	26.7	0.0	-.15[-.43,.06]	59.3
							Panel B Long-Short Equity						
Return	7.1	10.2	.03[-.35,.41]	23.3	3.5	-.13[-.55,.28]	19.2	0.0	-.20[-.66,.19]	30.8	0.0	-.39[-.76,-.05]	58.7
Alpha	9.1	11.8	.03[-.35,.36]	28.2	1.6	-.22[-.65,.15]	24.0	4.8	-.15[-.61,.24]	25.5	0.0	-.18[-.52,.10]	58.9
							Panel C Market Neutral						
Return	12.2	4.0	.03[-.35,.41]	15.7	0.0	-.11[-.32,.10]	24.2	6.7	-.10[-.37,.15]	18.6	0.3	-.04[-.28,.15]	53.1
Alpha	22.3	11.5	.03[-.35,.36]	14.2	4.3	-.01[-.33,.18]	20.4	13.2	-.09[-.34,.20]	18.9	11.9	-.04[-.27,.19]	57.4
							Panel D Managed Futures						
Return	6.8	6.8	-.02[-.34,.36]	5.8	13.7	.04[-.40,.50]	0.0	16.7	.16[-.21,.62]	1.5	0.0	-.00[-.33,.30]	31.1
Alpha	10.2	9.6	-.02[-.36,.37]	5.1	14.7	.11[-.32,.55]	1.2	13.1	.14[-.25,.54]	6.8	0.0	-.06[-.38,.25]	34.1
							Panel E Global Macro						
Return	0.0	13.2	.06[-.27,.42]	19.8	4.2	-.14[-.69,.23]	14.5	10.2	.02[-.55,.36]	28.3	0.0	-.19[-.54,.18]	46.2
Alpha	4.7	17.3	.08[-.25,.44]	19.8	0.0	-.12[-.68,.21]	15.8	9.0	-.08[-.62,.27]	21.7	5.2	-.07[-.42,-.24]	42.5

Table II

Return versus Alpha Predictability (Continued)

.

	Default Spread			Dividend Yield			VIX(Volatility)			Aggregate Flows			Joint
	Prop.(%)		Slope Coeff.	Prop.(%)		Slope Coeff.	Prop.(%)		Slope Coeff.	Prop.(%)		Slope Coeff.	Prop.(%)
Return	$\widehat{\pi}_R^-$	$\widehat{\pi}_R^+$	$\overline{b}_j[.25, .75]$	$\widehat{\pi}_R^-$	$\widehat{\pi}_R^+$	$\overline{b}_j[.25, .75]$	$\widehat{\pi}_R^-$	$\widehat{\pi}_R^+$	$\overline{b}_j[.25, .75]$	$\widehat{\pi}_R^-$	$\widehat{\pi}_R^+$	$\overline{b}_j[.25, .75]$	$\widehat{\pi}_R^{Joint}$
Alpha	$\widehat{\pi}_\alpha^-$	$\widehat{\pi}_\alpha^+$	$\overline{a}_j[.25, .75]$	$\widehat{\pi}_\alpha^-$	$\widehat{\pi}_\alpha^+$	$\overline{a}_j[.25, .75]$	$\widehat{\pi}_\alpha^-$	$\widehat{\pi}_\alpha^+$	$\overline{a}_j[.25, .75]$	$\widehat{\pi}_\alpha^-$	$\widehat{\pi}_\alpha^+$	$\overline{a}_j[.25, .75]$	$\widehat{\pi}_\alpha^{Joint}$
							Panel F Emerging Market						
Return	0.0	33.5	.43[-.03,1.14]	14.8	2.5	-.07[-.63,.30]	27.3	0.0	-.37[-1.08,.07]	37.9	0.0	-.54[-1.05,.16]	70.3
Alpha	0.0	37.9	.40[-.03,1.04]	20.9	2.4	-.19[-.70,.24]	37.9	3.0	-.45[-1.18,.07]	25.9	0.0	-.27[-.71,.04]	65.7
							Panel G Convertible Arbitrage						
Return	5.0	32.3	.12[-.07,.39]	32.6	2.7	-.11[-.34,.06]	24.9	17.9	-.06[-.39,.21]	44.7	0.0	-.17[-.46,-.05]	83.7
Alpha	9.0	28.0	.06[-.13,.34]	30.9	9.4	-.08[-.35,.14]	22.2	20.2	-.01[-.41,.30]	39.3	0.0	-.15[-.30,.01]	82.6
							Panel H Event-Driven						
Return	8.8	16.6	.05[-.20,.27]	37.8	2.6	-.12[-.46,.08]	39.0	0.0	-.26[-.62,.01]	34.4	0.0	-.19[-.42,-.03]	71.1
Alpha	18.0	8.3	-.03[-.27,.17]	28.5	0.0	-.11[-.43,.07]	30.8	1.0	-.14[-.47,.05]	22.2	0.0	-.09[-.34,.07]	56.1
							Panel I Fixed Income						
Return	10.2	20.7	.04[-.25,.26]	31.1	9.0	-.07[-.35,.12]	30.8	13.1	-.08[-.55,.20]	25.0	0.0	-.09[-.32,.04]	62.9
Alpha	18.0	19.5	.02[-.24,.20]	17.2	10.2	-.04[-.31,.17]	23.1	18.6	-.02[-.37,.29]	17.9	4.7	-.04[-.32,.09]	63.3
							Panel J Funds of Funds						
Return	0.0	15.1	.10[-.09,.26]	28.7	0.0	-.16[-.37,.04]	36.4	0.0	-.23[-.44,-.04]	65.1	0.0	-.29[-.52,-.14]	83.6
Alpha	3.0	19.7	.06[-.12,.22]	27.6	0.0	-.12[-.35,.06]	29.9	0.0	-.17[-.40,.00]	46.9	0.0	-.19[-.41,-.06]	79.2
							Panel K Multi-Strategy						
Return	1.6	9.9	.07[-.24,.42]	21.1	11.9	-.06[-.39,.42]	3.4	13.7	.11[-.26,.51]	7.5	0.0	-.10[-.38,.15]	40.5
Alpha	4.3	14.1	.07[-.26,.41]	11.5	9.6	-.03[-.40,.38]	3.8	7.9	.08[-.28,.46]	19.0	0.0	-.13[-.44,.10]	44.3

Table III
The Economic Value of Predictability

We compare the performance of the unconditional strategy with the one produced by different conditional strategies. The unconditional strategy selects funds with the highest unconditional t-statistic, while the conditional strategies (single-, multi-predictor, and combination) selects funds with the highest conditional t-statistic. While the single-predictor strategies use one of the four predictors (default spread, dividend yield, VIX, and aggregate flows) to forecast returns, the multi-predictor strategy uses all predictors simultaneously. Finally, the combination strategy averages across the single-predictor conditional t-statistics. All portfolios are formed at the end of the year and rebalanced annually. The initial formation date is on December 31, 1996, and the final one on December 31, 2006. For comparison purposes, we also report the performance of hedge fund value-weighted (VW) and equally-weighted (EW) indices. In Panel A, we report the (annualized) Fung-Hsieh alpha, $\widehat{\alpha}$, residual standard deviation, $\widehat{\sigma}_{res}$, and Information Ratio, IR, the (annualized) excess mean, $\widehat{\mu}$, standard deviation, $\widehat{\sigma}_{tot}$, and Sharpe Ratio, SR, as well as the 1%- and 5%-Value-at-Risk, VaR. In Panel B, we compute the median and the 25-75% quantiles for the number of funds chosen each year, as well as the annual turnover. We also display the estimated first- and second-lag Fung-Hsieh residual autocorrelation and the average weights associated with the three investment categories in which the strategies invest most (LS, MF, and ED denote Long-Short, Managed Futures, and Event Driven, respectively). Panel C reports the portfolio sensitivity to the Fung-Hsieh risk factors along with the model explanatory power, R^2. Figures in parentheses denote the bootstrap p-values under the null hypothesis that the parameter is equal or inferior to the one associated with the unconditional strategy (one-sided test).

Panel A Out-of-Sample Performance (Jan. 1997-Dec. 2007)

	Fung-Hsieh Alpha (Ann.)			Excess Return (Ann.)			VaR	
	$\widehat{\alpha}$	$\widehat{\sigma}_{res}$	IR=$\widehat{\alpha}/\widehat{\sigma}_{res}$	$\widehat{\mu}$	$\widehat{\sigma}_{tot}$	SR=$\widehat{\mu}/\widehat{\sigma}_{tot}$	1%	5%
Unconditional	5.8%	2.4%	2.4	7.3%	4.2%	1.8	-1.5%	-0.9%
Single-Predic.								
Default Spread	7.8(.06)	3.9	2.0(.89)	9.2(.07)	5.7	1.6(.81)	-2.8	-1.0
Dividend Yield	6.5(.17)	3.5	1.8(.98)	8.5(.05)	5.7	1.5(.92)	-1.9	-1.2
VIX (Volatility)	6.2(.29)	3.5	1.8(.97)	7.8(.16)	5.0	1.5(.77)	-3.6	-1.5
Aggregate Flows	5.6(.62)	2.8	2.0(.91)	7.3(.49)	4.5	1.6(.82)	-2.6	-1.0
Multi-Predic.	5.3(.69)	3.4	1.5(.98)	6.6(.73)	4.6	1.4(.90)	-2.9	-1.3
Combination	7.0(.00)	2.6	2.7(.03)	8.5(.00)	4.5	1.9(.05)	-1.4	-0.7
Index (VW)	3.7(.96)	3.8	1.0(1.0)	5.5(.98)	5.3	1.0(1.0)	-2.6	-2.1
Index (EW)	4.1(.95)	3.1	1.3(1.0)	5.6(.98)	5.4	1.0(1.0)	-2.6	-2.0

Table III

The Economic Value of Predictability (Continued)

Panel B Portfolio Characteristics

	Nb. funds	Turn. (Ann)	Residual Autocorr. lag 1	Residual Autocorr. lag 2	Investment Categories LS	Investment Categories MF	Investment Categories ED
Unconditional	69[63,71]	59%[51,66]	.14	-.05	13.1%	6.2%	5.7%
Single-Predic.							
Default Spread	68[61,71]	80[70,85]	.22(.04)	.03(.13)	13.7	7.2	6.3
Dividend Yield	68[64,72]	76[61,77]	.21(.04)	.02(.09)	14.5	5.7	6.7
VIX (Volatility)	69[62,71]	68[64,78]	.03(.99)	.00(.06)	15.0	7.7	5.9
Aggregate Flows	70[61,72]	71[59,75]	.24(.01)	-.04(.79)	12.3	7.0	7.0
Multi-Predic.	68[59,72]	86[84,92]	.10(.78)	.00(.06)	11.9	10.5	6.3
Combination	69[62,71]	66[58,71]	.18(.13)	-.04(.52)	11.8	6.2	7.2
Index (VW)	808[606,947]	27[25,29]	.04(.97)	-.06(.37)	15.2	10.3	5.9
Index (EW)	808[606,947]	30[28,32]	.06(.92)	-.04(.27)	20.7	12.0	5.2

Panel C Sensitivity to the Fung-Hsieh Risk Factors

	Market	Size	Term	Default	T.Bond	T.Cur.	T.Com.	R^2
Unconditional	.15	.08	.10	.07	-.02	.00	.01	50.6%
Single-Predic.								
Default Spread	.19(.03)	.11(.07)	.08(.73)	.10(.40)	-.01(.11)	.00(.76)	.00(.90)	38.0
Dividend Yield	.20(.00)	.14(.00)	.18(.01)	.11(.24)	-.02(.74)	.01(.35)	.01(.38)	49.9
VIX (Volatility)	.14(.46)	.06(.93)	.11(.43)	.18(.03)	-.01(.27)	.01(.17)	.02(.00)	34.0
Aggregate Flows	.14(.68)	.07(.77)	.09(.65)	.25(.00)	-.01(.23)	.01(.42)	.00(.68)	47.8
Multi-Predic.								
	.09(.99)	.09(.36)	.09(.56)	.23(.02)	-.01(.10)	.01(.26)	.01(.36)	28.0
Combination	.14(.26)	.09(.13)	.10(.62)	.12(.04)	-.01(.38)	.00(.65)	.00(.81)	51.6
Index (VW)	.17(.07)	.10(.19)	.16(.10)	.25(.01)	.00(.00)	.01(.03)	.02(.02)	41.2
Index (EW)	.22(.00)	.15(.00)	.12(.25)	.17(.05)	.01(.00)	.02(.00)	.02(.04)	60.8

Table IV

Predictor Value and Performance
Single-Predictor Strategies

For each predictor (default spread, dividend yield, VIX, and aggregate flows), we sort its values observed on the 11 portfolio formation dates (from December 31, 1996 to December 31, 2006) into quintiles, and focus on the bottom and top quintiles. In Panel A, we report the characteristics (observed on the formation date) of each single-predictor strategy for the two quintiles (bottom and top). We compute the predictor value, $z_{j,t}$, the proportion of funds that are common to those held in the unconditional and slope portfolios (these are denoted by $\%_{uncond.}$ and $\%_{slope}$, respectively), and the average slope t-statistic, $t(\widehat{b}_j)$, across the funds selected by the strategy. Each value, $z_{j,t}$, is standardized, i.e., a value of one means that $z_{j,t}$ is one standard deviation higher than its normal level. The estimated slope coefficient of each fund is multiplied by $sign(z_{j,t})$ to guarantee that it is positive only when the fund has the right exposure to the predictor (i.e., when $sign(\widehat{b}_{i,j}) = sign(z_{j,t})$). In Panel B, we report the subsequent annual performance of each strategy following the formation date for the two quintiles (bottom and top). We report the (annualized) Fung-Hsieh alpha, $\widehat{\alpha}$, and Information Ratio, IR, as well as the (annualized) excess mean, $\widehat{\mu}$, and Sharpe Ratio, SR.

Panel A Portfolio Characteristics (Formation Dates)

	Bottom Quintile				Top Quintile			
	$z_{j,t}$	$\%_{uncond.}$	$\%_{slope}$	$t(\widehat{b}_j)$	$z_{j,t}$	$\%_{uncond.}$	$\%_{slope}$	$t(\widehat{b}_j)$
Default Spread	-1.65	58.1%	43.4%	0.96	4.81	28.2%	77.2%	2.03
Dividend Yield	-1.77	40.7	57.3	2.44	1.62	52.7	32.7	0.72
VIX (Volatility)	-1.16	51.3	46.0	1.84	2.04	44.0	56.7	1.04
Aggregate Flows	-1.26	64.0	27.3	0.87	1.29	56.9	37.6	0.36

Panel B Out-of-Sample Performance

	Bottom Quintile				Top Quintile			
	$\widehat{\alpha}$	IR	$\widehat{\mu}$	SR	$\widehat{\alpha}$	IR	$\widehat{\mu}$	SR
Default Spread	6.0%	1.9	7.5%	2.4	14.8%	2.2	14.4%	2.1
Dividend Yield	9.6	2.0	12.9	2.7	5.0	1.3	8.7	2.4
VIX (Volatility)	6.9	2.6	10.1	3.9	1.1	0.2	5.2	1.0
Aggregate Flows	5.6	1.5	6.7	1.8	2.3	1.1	7.6	3.4

Table V

Predictor Value and Performance
Combination Strategy

For each predictor (default spread, dividend yield, VIX, and aggregate flows), we sort its values observed on the 11 portfolio formation dates (from December end 1996 to December end 2006) into quintiles, and focus on the bottom and top quintiles. For each predictor and each quintile (bottom and top), Panel A displays the characteristics (observed on the formation date) and the performance (measured over the subsequent year) of the combination long portfolio that only invests in funds that are included in the combination strategy but not in the unconditional strategy. We compute the predictor value, $z_{j,t}$, and the proportion of funds in the long portfolio that are common to those held by the slope portfolio, $\%_{slope}$. Each value, $z_{j,t}$, is standardized, i.e., a value of one means that the predictor value is one standard deviation higher than its normal level. We also report the difference in Fung-Hsieh alpha and mean between the long portfolio and the unconditional strategy, denoted by $\widehat{\alpha}$-$\widehat{\alpha}_u$ and $\widehat{\mu}$-$\widehat{\mu}_u$, respectively. Panel B repeats this analysis for the combination short portfolio that only invests in funds that are included in the unconditional strategy but not in the combination strategy.

Panel A Long Portfolio

	Bottom Quintile				Top Quintile			
	$z_{j,t}$	$\%_{slope}$	$\widehat{\alpha}-\widehat{\alpha}_u$	$\widehat{\mu}-\widehat{\mu}_u$	$z_{j,t}$	$\%_{slope}$	$\widehat{\alpha}-\widehat{\alpha}_u$	$\widehat{\mu}-\widehat{\mu}_u$
Default Spread	-1.65	8.1%	5.3%	5.2%	4.81	10.1%	4.7%	4.6%
Dividend Yield	-1.77	11.2	3.2	3.7	1.62	6.0	1.9	3.5
VIX (Volatility)	-1.16	14.0	2.4	4.7	2.04	9.1	3.2	3.7
Aggregate Flows	-1.26	2.0	1.5	2.2	1.29	8.3	0.0	3.1

Panel B Short Portfolio

	Bottom Quintile				Top Quintile			
	$z_{j,t}$	$\%_{slope}$	$\widehat{\alpha}-\widehat{\alpha}_u$	$\widehat{\mu}-\widehat{\mu}_u$	$z_{j,t}$	$\%_{slope}$	$\widehat{\alpha}-\widehat{\alpha}_u$	$\widehat{\mu}-\widehat{\mu}_u$
Default Spread	-1.65	0.0%	-5.5%	-3.7%	4.81	0.0%	-2.7%	-3.7%
Dividend Yield	-1.77	0.0	-2.0	0.3	1.62	0.0	0.0	-1.3
VIX (Volatility)	-1.16	0.0	-1.5	-1.1	2.04	0.0	-2.0	0.3
Aggregate Flows	-1.26	0.0	-0.5	-1.9	1.29	0.0	2.2	2.9

Table VI

Impact of the 2008 Financial Crisis

In Panel A, we compare the performance of the unconditional strategy with the one obtained by different conditional strategies (single-, multi-predictor, and combination) after including the 2008 financial crisis. We report the (annualized) Fung-Hsieh alpha, $\widehat{\alpha}$, residual standard deviation, $\widehat{\sigma}_{res}$, and Information Ratio, IR, the (annualized) excess mean, $\widehat{\mu}$, standard deviation, $\widehat{\sigma}_{tot}$, and Sharpe Ratio, SR, as well as the 2008 cumulative returns during the periods January-September and October-December. Figures in parentheses denote the bootstrap p-values under the null hypothesis that the parameter is equal or inferior to the one associated with the unconditional strategy (one-sided test). In Panel B, we examine, for each single-predictor strategy, the stability of the estimated slope coefficient in 2008. We first report the predictor value, z_j, as well as the proportion of funds that are common to those held in the unconditional and slope portfolios, $\%_{uncond.}$ and $\%_{slope}$, on the final formation date in December 31, 2007. Each value, z_j, is standardized, i.e., a value of one means that z_j is one standard deviation higher than its normal level. Then, we compute two hit ratios that determine the proportion of months in 2008 when: 1) the estimated portfolio predictable component is positive ($\widehat{b}_{j,t} z_{j,t} > 0$); 2) the predictor value has the same sign as the one observed on December 31, 2007 ($z_j z_{j,t} > 0$). For comparative purposes, we report the same statistics computed during previous years (excluding 2008) following an extreme predictor value (we use the values included in either the bottom or top quintile in Table IV, depending on the sign of z_j). In Panel C, we measure the performance of the unconditional and unconditional strategies between January 1997 and December 2008, assuming that the investor can rebalance his portfolio at a monthly frequency in 2008. The Δ reports the difference in performance with the baseline strategies in Panel A.

Panel A Out-of-Sample Performance (Jan. 1997-Dec. 2008)

	Fung-Hsieh Alpha (Ann.)			Excess Return (Ann.)			Cum. Return 2008	
	$\widehat{\alpha}$	$\widehat{\sigma}_{res}$	IR	$\widehat{\mu}$	$\widehat{\sigma}_{tot}$	SR	Jan-Sep	Oct-Dec
Unconditional	4.1%	3.4%	1.2	4.9%	5.1%	0.9	-5.4%	-14.1%
Single-Predictor								
Default Spread	6.2(.02)	4.2	1.5(.11)	6.8(.03)	6.3	1.1(.14)	-4.9	-12.3
Dividend Yield	5.8(.03)	3.8	1.5(.07)	7.0(.00)	5.9	1.2(.07)	0.1	-8.6
VIX (Volatility)	6.3(.01)	3.8	1.7(.08)	7.2(.01)	5.1	1.4(.02)	4.7	-3.1
Aggregate Flows	4.2(.41)	3.4	1.2(.44)	4.9(.49)	5.3	0.9(.60)	-5.1	-14.1
Multi-Predictor	4.9(.22)	3.6	1.3(.32)	5.6(.26)	4.7	1.2(.14)	0.7	-4.7
Combination	6.0(.00)	3.1	1.9(.00)	6.7(.00)	5.1	1.3(.00)	0.2	-11.2

Table VI

Impact of the 2008 Financial Crisis (Continued)

Panel B Stability of the Slope Coefficient

	December 07 (Formation date)			Jan.-Dec. 2008 Hit Ratios (> 0)		Jan. 1997-Dec. 2007 Hit Ratios (> 0)	
	z_j	$\%_{uncond.}$	$\%_{slope}$	$\widehat{b}_{j,t} \cdot z_{j,t}$	$z_j \cdot z_{j,t}$	$\widehat{b}_{j,t} \cdot z_{j,t}$	$z_j \cdot z_{j,t}$
Single-Predictor							
Default Spread	3.17	24.0%	62.7%	25.0%	100%	87.5%	100%
Dividend Yield	3.36	21.3	73.3	16.6	100	83.3	100
Volatility (VIX)	2.46	24.0	76.0	83.3	100	70.8	95.8
Aggregate Flows	-1.33	54.7	48.0	50.0	75.0	70.8	66.6

Panel C Impact of Liquidity Constraints

	Fung-Hsieh Alpha (Ann.)				Excess Return (Ann.)				Cum. Return 2008	
	$\widehat{\alpha}$	Δ	IR	Δ	$\widehat{\mu}$	Δ	SR	Δ	Jan-Sep	Oct-Dec
Unconditional	5.6%	1.5	1.8	0.6(.00)	6.4%	1.5	1.5	0.6(.01)	-1.6%	-0.6%
Single-Predictor										
Default Spread	8.2	2.0	1.6	0.1(.28)	8.6	1.8	1.4	0.3(.06)	-4.0	7.6
Dividend Yield	7.3	1.5	1.5	0.0(.56)	8.3	1.3	1.3	0.1(.17)	-2.6	9.7
VIX (Volatility)	6.9	0.6	1.6	-0.1(.71)	7.7	0.5	1.5	0.1(.30)	-2.0	6.1
Aggregate Flows	6.6	2.4	1.6	0.4(.06)	7.0	2.1	1.4	0.5(.01)	-0.7	5.9
Multi-Predictor	4.7	-0.2	1.2	-0.1(.89)	5.4	0.2	1.1	-0.1(.79)	-1.6	-4.6
Combination	7.7	1.7	1.9	0.0(.51)	8.2	1.5	1.6	0.3(.06)	-0.5	7.4

Table VII
Sensitivity Analysis

We examine whether performance is sensitive to changes in the baseline specification. We consider the following cases: the maximum number of funds in the portfolio is either equal to 50 or unlimited (baseline case: 75 funds); there is no AuM cutoff (baseline case: small funds are discarded); there is a -25% monthly return penalty when a fund has a missing observation (baseline case: 0% penalty); we assume that returns over the previous month are not available in the hedge fund databases (baseline case: 0 month); we impose a three-month notice period prior to withdrawal (baseline case: 0 months); we use the t-statistic of the fund conditional alphas to select funds (baseline case: conditional mean). In each case, we report the (annualized) Fung-Hsieh alpha, $\widehat{\alpha}$, Information Ratio, IR, excess mean, $\widehat{\mu}$, and Sharpe Ratio, SR, during the periods January 1997-December 2007 and January 1997-December 2008. Figures in parentheses denote the bootstrap p-values under the null hypothesis that the difference in IR (SR) between the combination and the unconditional strategies is zero or less (one-sided test).

		Jan. 1997-Dec. 2007				Jan. 1997-Dec. 2008			
		$\widehat{\alpha}$	IR	$\widehat{\mu}$	SR	$\widehat{\alpha}$	IR	$\widehat{\mu}$	SR
Portfolio Size:	Unconditional	5.9	2.7	7.1	1.9	4.5	1.6	5.2	1.2
50 Funds	Combination	7.3	3.0(.08)	8.6	2.0(.08)	6.3	2.2(.02)	6.9	1.5(.01)
Portfolio Size:	Unconditional	6.1	2.4	7.6	1.7	4.4	1.3	5.2	1.0
Unlimited	Combination	7.0	2.5(.14)	8.5	1.8(.08)	5.8	1.8(.00)	6.6	1.2(.00)
AUM Cutoff:	Unconditional	6.7	2.9	8.0	1.9	5.9	2.3	6.5	1.5
No Cutoff	Combination	7.5	3.2(.03)	9.0	2.0(.09)	6.7	2.6(.04)	7.4	1.6(.13)
Missing Return:	Unconditional	3.3	1.1	4.6	1.1	1.5	0.4	2.2	0.4
-25% Penalty	Combination	4.6	1.5(.01)	5.8	1.3(.01)	3.4	1.0(.00)	4.0	0.8(.00)
Data Availability:	Unconditional	5.4	2.1	6.9	1.6	3.7	1.1	4.5	0.9
1 Month Lag	Combination	6.3	2.4(.07)	8.0	1.7(.09)	5.2	1.6(.00)	6.0	1.1(.00)
Notice Period:	Unconditional	5.0	2.1	6.3	1.6	3.4	1.1	4.1	0.8
3 Months	Combination	5.4	2.2(.24)	6.3	1.8(.10)	4.4	1.6(.01)	4.8	1.2(.00)
Alpha	Unconditional	6.4	2.5	7.7	1.9	5.2	1.9	5.9	1.4
Predictability	Combination	7.7	2.9(.07)	8.9	2.0(.27)	6.7	2.2(.06)	7.1	1.5(.19)

Figure 1

The Trade-off between Unconditional and Predictable Performance

We plot the relation between the predictor value and the conditional t-statistic of the unconditional and slope portfolios in Panels A and B, respectively. The unconditional (slope) portfolio holds the top decile of funds with the highest unconditional (slope) t-statistic. Each graph is constructed from our hedge fund dataset (to be presented), using the default spread as a single predictor. For each fund i included in the unconditional (slope) portfolio in month t, we use past data to compute $t(\widehat{\mu}_i)$, $\widehat{var}(\widehat{\mu}_i)$, $t(\widehat{b}_{i,j})$, $\widehat{var}(\widehat{b}_{i,j})$, and $\widehat{var}(\widehat{\mu}_{i,t})$. Then, we average these estimates across funds and months, and plug these averages into Equation (7) to obtain the conditional t-statistic for each portfolio. In Panel C, we illustrate the investment process of the conditional strategy that holds the top decile of funds with the highest conditional t-statistic.

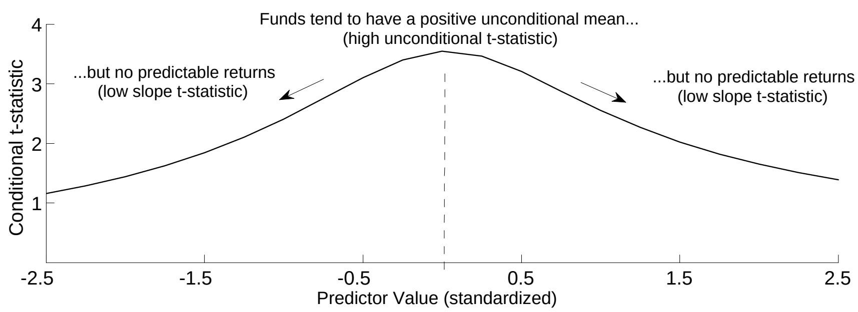

(A) Unconditional Portfolio

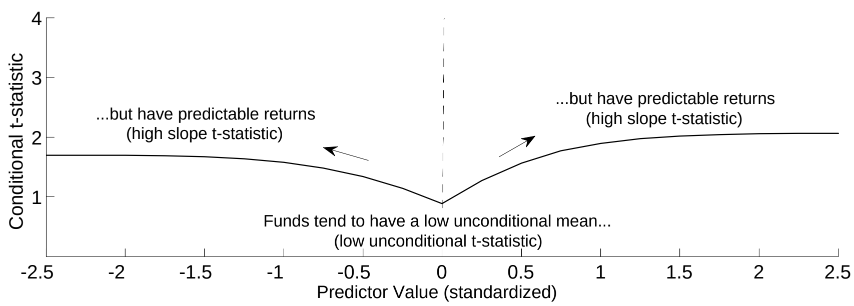

(B) Slope Portfolio

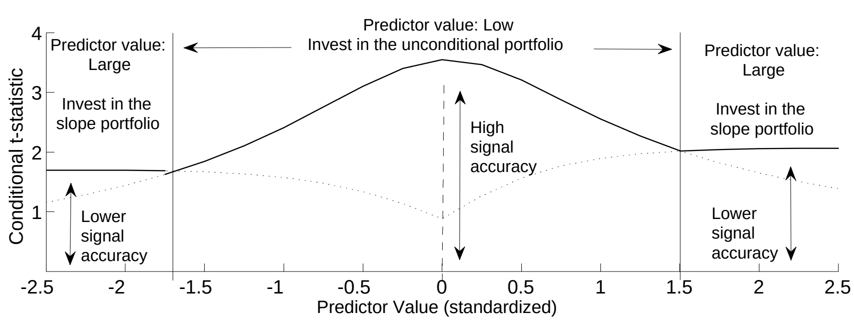

(C) Conditional Strategy

Figure 2

Evolution of the Predictive Variables

We report monthly values of the four predictors between January 1994 and December 2008 (e.g., the January observation corresponds to the predictor value observed in December that is used to forecast returns in January). The default spread (Panel A) is the yield differential between Moodys BAA-rated and AAA-rated bonds. The dividend yield (Panel B) is the total cash dividends on the value-weighted CRSP index over the previous 12 months divided by the current level of the index. The VIX (Panel C) correspond to the CBOE volatility index. Aggregate flows (Panel D) is calculated as the value-weighted monthly percentage in- and outflows into the hedge funds contained in our database.

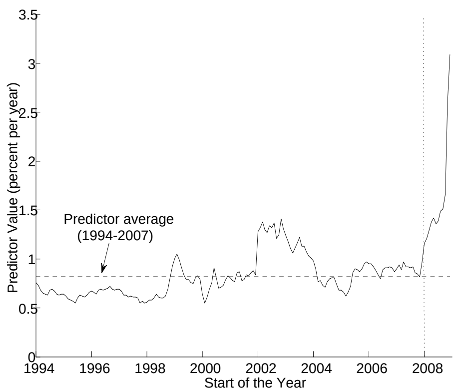

(A) Default Spread

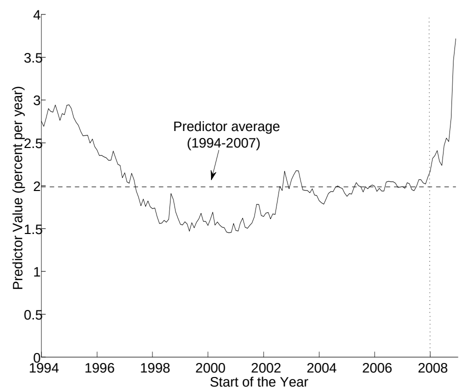

(B) Dividend Yield

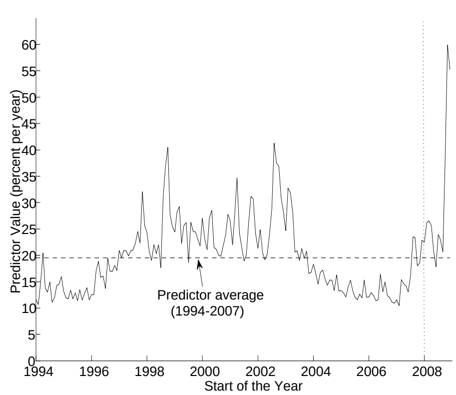

(C) VIX (Volatility)

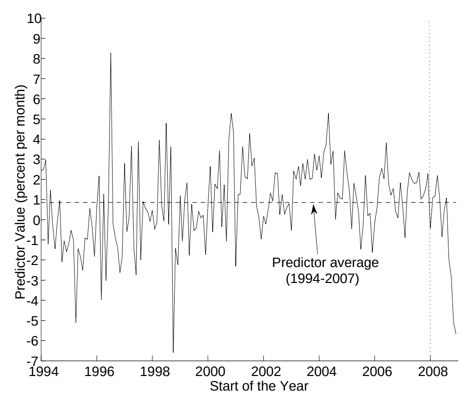

(D) Aggregate Flows

Figure 3

Time Variation in Cumulative Wealth

We plot the evolution one dollar invested in the combination strategy and in the unconditional portfolio from January 1997 until December 2007.

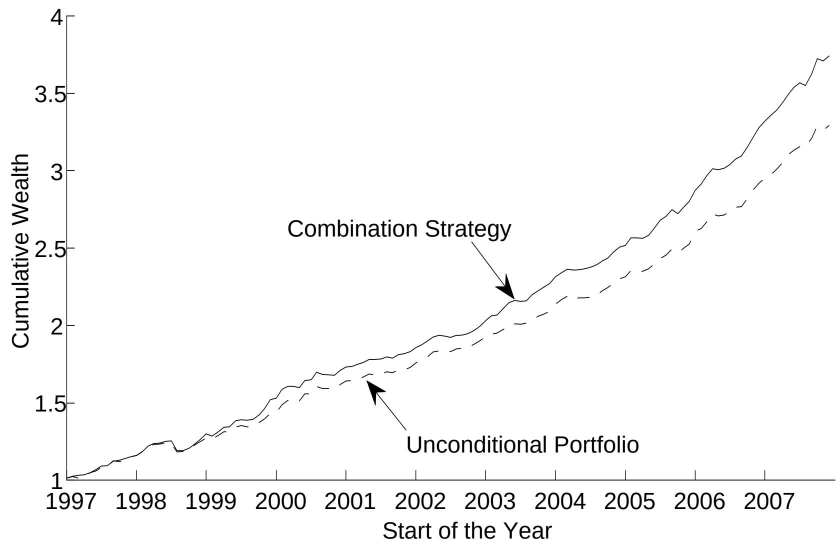

Figure 4

Impact of Changing Predictor Values on Performance Single-Predictor Strategies

For each predictor, we examine the relation between its value during the portfolio formation date (December end) and the change in performance of the single-predictor strategy over the subsequent year. The results for the default spread, the dividend yield, the VIX, and aggregate flows are displayed in Panels A-D. The predictor value is standardized, i.e., a value of one means that its value is one standard deviation higher than its normal level. Performance is measured as the difference in Information Ratio between the single-predictor strategy and the unconditional portfolio, computed using an expanding window that includes the return data up to the end of the year. Specifically, the initial observation in 1997 corresponds to the predictor value observed in December 1996 along with the portfolio performance measured from January 1997 to December 1997. The final observation in 2007 corresponds to the predictor value in December 2006 along with the portfolio performance measured from January 1997 to December 2007.

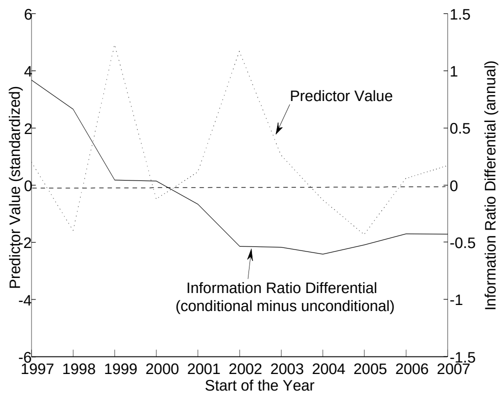

(A) Default Spread

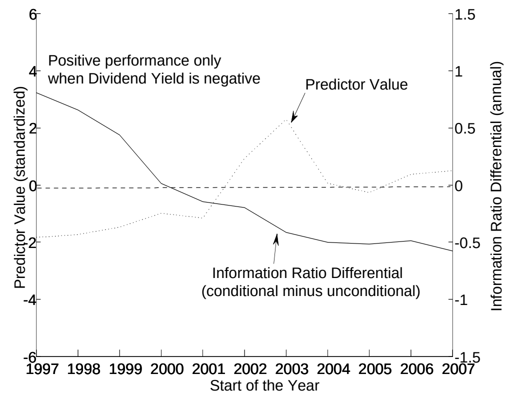

(B) Dividend Yield

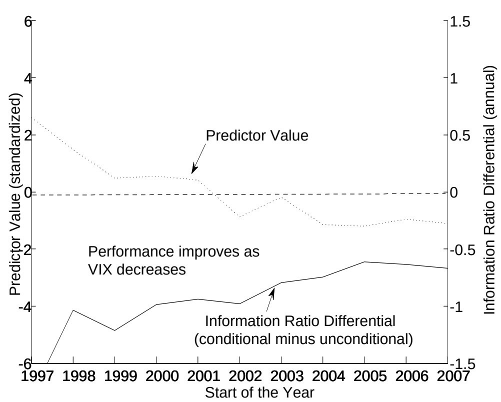

(C) VIX (Volatility)

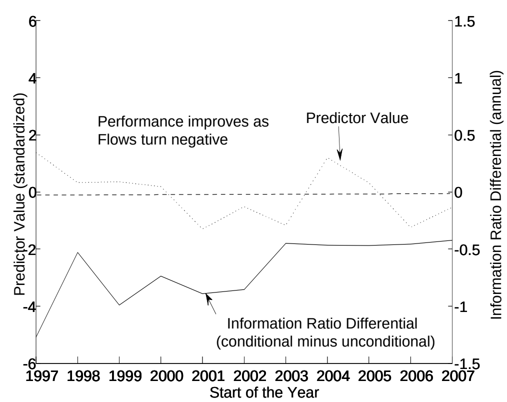

(D) Aggregate Flows

Figure 5

Commonality with the Unconditional Portfolio

We show the evolution of the proportion of funds chosen by each conditional strategy that are common to those included in the unconditional portfolio between January 1997 and December 2007. Panels A to D show the results for the single-predictor strategies (default spread, dividend yield, VIX, and aggregate flows). Panels E and F display the results for the multi-predictor and the combination strategies, respectively.

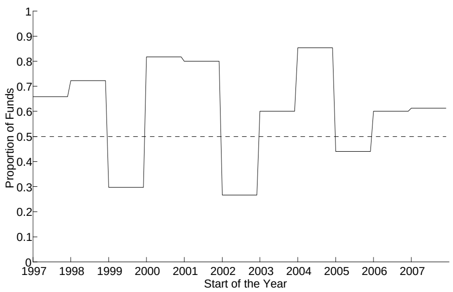

(A) Default Spread

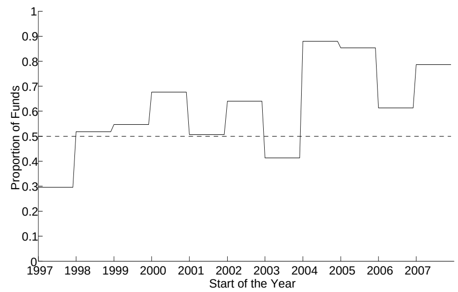

(B) Dividend Yield

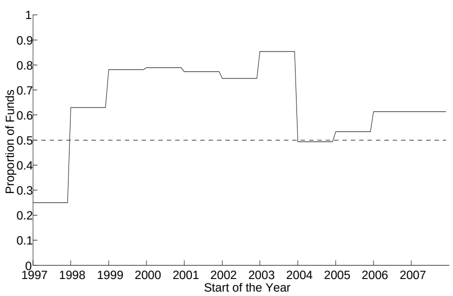

(C) VIX (Volatility)

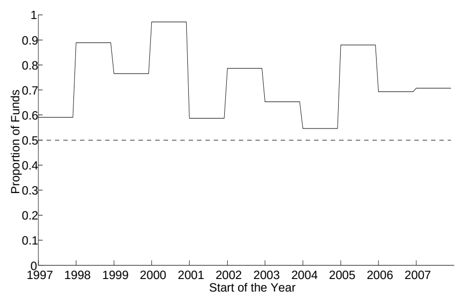

(D) Aggregate Flows

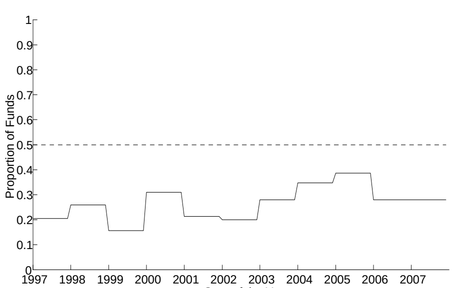

(E) Multi-Predictor

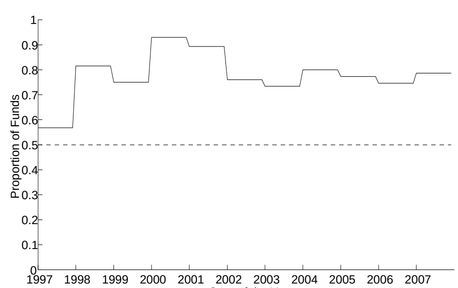

(F) Combination Strategy

Figure 6

Allocation of the VIX Strategy across Investment Styles

Panel A shows the standardized value of the VIX (Volatility) on each formation date (i.e., one means that the VIX is one standard deviation higher than its normal level). The initial observation is the value observed in December 1996, while the final one is observed in December 2007. The monthly weights invested across investment styles by the unconditional and the VIX strategies are shown in Panels B and C, respectively. To improve readability, we discard any investment category for which the associated weight is below 5% during the entire investment period. In addition, we scale up the weights by discarding the proportion invested the category "Other" (this proportion is around 50%, and is very similar across strategies).

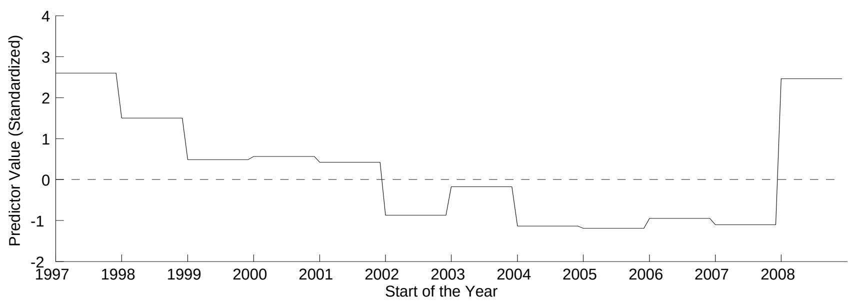

(A) Predictor Value (Volatility)

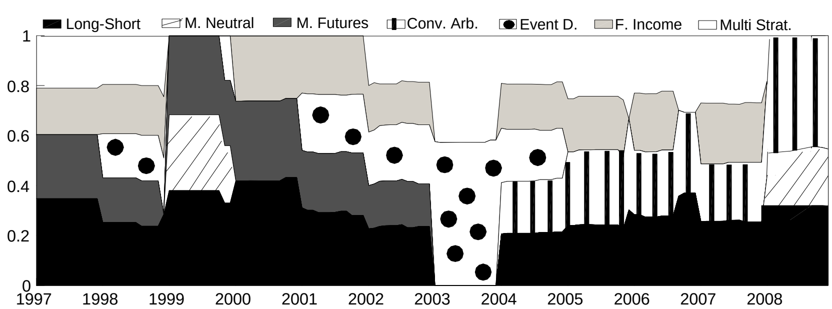

(B) Portfolio Composition: Unconditional Portfolio

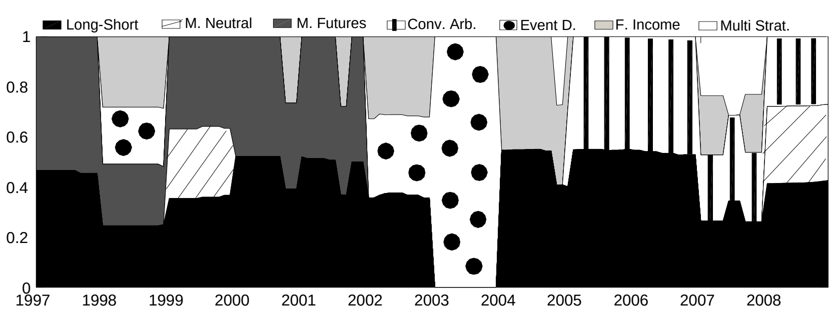

(C) Portfolio Composition: Conditional Strategy (VIX)