The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXVIII, NO. 2 • APRIL 2013

On the High-Frequency Dynamics of Hedge Fund Risk Exposures

ANDREW J. PATTON and TARUN RAMADORAI*

ABSTRACT

We propose a new method to model hedge fund risk exposures using relatively high-frequency conditioning variables. In a large sample of funds, we find substantial evidence that hedge fund risk exposures vary across and within months, and that capturing within-month variation is more important for hedge funds than for mutual funds. We consider different within-month functional forms, and uncover patterns such as day-of-the-month variation in risk exposures. We also find that changes in portfolio allocations, rather than in the risk exposures of the underlying assets, are the main drivers of hedge funds' risk exposure variation.

AN IMPORTANT FEATURE OF hedge funds is the speed with which they alter their investments in response to changing market conditions. Static analyses of hedge funds' risk exposures are likely to miss these rapid changes in their strategies or leverage ratios, and thus several new approaches have been proposed to model the dynamics of these risk exposures.[1] One factor that must be taken into consideration is the high frequency (often daily or even higher) with which hedge fund risk exposures change.[2] However, the approaches proposed thus far to model dynamic risk exposures track such changes only at the monthly frequency, as this is the reporting frequency for performance data in all of the main hedge fund databases.

We propose a new method to surmount this obstacle, and to better understand the high-frequency dynamics of hedge fund risk exposures. The starting point

*Patton is with Duke University and the Oxford-Man Institute of Quantitative Finance. Ramadorai is with Saïd Business School, Oxford University, Oxford-Man Institute of Quantitative Finance, and CEPR. Acting Editor: Jennifer Conrad. We thank Alexander Taylor and Sushant Vale for dedicated research assistance and Nick Bollen, Michael Brandt, Mardi Dungey, Jean-David Fermanian, Robert Kosowski, Olivier Scaillet, Kevin Sheppard, Melvyn Teo, and seminar participants at Fuqua School of Business, the Oxford-Man Institute Hedge Fund Conference, the CREST-HEC Hedge Fund Conference, the 2010 SoFiE annual conference, Lancaster University, the University of Tasmania, and the 2011 Western Finance Association conference for useful comments.

[1] The literature on modeling hedge fund returns using static models is extensive. A partial list includes Fung and Hsieh (1997, 2004a,b), Ackermann, McEnally, and Ravenscraft (1999), Liang (1999), Agarwal and Naik (2004), Kosowski, Naik, and Teo (2007), Chen and Liang (2007), Fung et al. (2008), Patton (2009), Jagannathan, Malakhov, and Novikov (2010), and Agarwal et al. (2011).

[2] See "Wall Street's New Race Toward Danger," *Barron's*, March 8, 2010 and "Traders Piqued by the Picosecond, but Physics Intervenes," *Wall Street Journal*, March 10, 2010.

for our approach is the widely used Ferson and Schadt (1996) model, which we extend to employ higher-frequency conditioning information. To circumvent the lack of high-frequency data on hedge fund performance, we posit a daily factor model for hedge fund returns and then aggregate this up to the monthly frequency for estimation. We demonstrate that the method is able to accurately track the dynamics of daily variation in hedge fund risk exposures by testing it on daily indexes of hedge fund returns, and by a simulation study.

The higher-frequency version of Ferson and Schadt (1996) in which daily risk exposures evolve as a linear function of observable instruments (which we dub the "linear model") is the first of three economically motivated functional forms for high-frequency risk exposures that we consider. The second model we consider allows for intra-month seasonalities in risk exposures (the "day-of-the-month model"), and the third model allows risk exposures to vary abruptly when observable instruments hit prespecified threshold values (the "threshold model"). By allowing for a variety of economically plausible ways in which risk exposures may evolve, we attempt to mitigate the inevitable loss of information that arises when using monthly returns to infer intra-monthly dynamics. More importantly perhaps, the results from these different specifications allow us to better understand hedge fund behavior during nonreporting intervals that have thus far been impervious to scrutiny.

We implement these dynamic risk exposure models on a cross-section of 14,194 individual hedge funds and funds-of-funds over the period 1994–2009, and find that they perform very well in explaining hedge fund returns. In particular, the models generate adjusted R^2 statistics that are a substantial improvement over a static-parameter benchmark model: the average adjusted R^2 for our linear specification is 49% higher than the corresponding average for the Fung–Hsieh benchmark model. We also find that including higher-frequency conditioning information substantially improves the performance of our model: the percentage of hedge funds for which we find statistically significant factor exposure variation nearly doubles, from 12% to 22%, when we include daily information as well as monthly information in our estimated specifications. In contrast, when we estimate our model on a set of 32,913 equity and bond mutual funds, adding daily information to the monthly information set leaves the percentage of funds for which we find statistically significant factor exposure variation virtually unchanged. In short, there is significant daily variation in hedge fund risk exposures, and accounting for this daily variation is necessary to characterize hedge fund behavior accurately whereas daily risk exposure variation does not seem to be as important for mutual funds, a fact that is perhaps unsurprising given short-sales restrictions and other constraints on mutual fund portfolio alterations.

The models that we propose provide new and valuable insights into hedge fund behavior at high frequencies. One particularly interesting finding from the day-of-the-month model is that there are significant intra-month seasonalities in hedge fund risk exposures. In particular, we find that hedge fund risk exposures are relatively high at the beginning of the month and decline steadily as the month progresses, reaching their lowest point at the end of the month just to the date at which hedge funds report returns to databases. There are several possible explanations for this phenomenon. One innocuous explanation is that this pattern reflects regular expirations of short-lived derivative positions held by funds. Another less innocuous explanation is that this pattern constitutes evidence of intra-month window-dressing by hedge funds. This explanation links our result to the extensive literature on mutual fund and pension fund window-dressing pioneered by Lakonishok et al. (1991), as well as the growing body of literature by authors such as Bollen and Krepely-Pool (2009) and Agarwal, Daniel, and Naik (2011) on unusual monthly patterns in hedge fund returns.

The threshold model also yields useful insights, one of which is that hedge funds have a tendency to abruptly cut positions in response to significant market events. For example, when market returns fall or when illiquidity rises significantly within the month, hedge funds significantly cut exposures to small stocks. This evidence is akin to that provided at lower frequencies by Brunnermeier and Nagel (2004), who document that hedge funds "rode" the technology bubble, and fits the description of destabilizing rational speculation provided in DeLong et al. (1990). Furthermore, when S&P 500 volatility rises significantly within the month, we find evidence that hedge funds cut back their positions across all risky assets, appearing to retreat towards cash at such times. This provides useful evidence in favor of Ferson and Schadt's (1996) conjecture that the enhanced mutual fund alpha that they detect using a time-varying beta model is on account of managers adjusting risk exposures in line with movements in aggregate volatility, and the more recent extension by Lo (2008), who decomposes fund manager performance into a "passive" component and an "active" component that arises from the correlation between changing portfolio weights and returns. Indeed, as we describe below, we also find improved performance from our time-varying beta model relative to the alpha obtained from a static factor model.

In the discussion thus far, we have tended to interpret evidence of time-varying risk exposures in terms of fund managers actively shifting portfolio allocations. Of course, it is possible that time variation in underlying asset betas could result in changes in fund risk exposures even when fund managers pursue passive buy-and-hold strategies. To evaluate their relative magnitudes, we posit a simple decomposition of fund beta variation into weight variation, asset beta variation, and weight-beta covariation. We then estimate this decomposition using matched 13-F data on all long-short equity hedge funds in our sample.[3] Using these data, we find that on average from 1989 to 2010, weight variation accounts for 73% of total fund beta variation and asset beta variation constitutes 17%, with covariances accounting for the remainder. During the recent financial crisis (2007–2010), weight variation accounts for an increased

[3] The quarterly 13-F filings data track long equity positions of institutional investment managers, and thus this analysis is naturally restricted to the subset of long-short equity fund managers in our data. This analysis complements our high-frequency analysis, providing us with additional information on the underlying causes for movements in hedge fund risk exposures. We thank an anonymous referee for suggesting that we pursue this strategy.

share (84% of the total), with the share of pure asset beta variation practically unchanged. In sum, the evidence indicates that the primary source of hedge funds' dynamic risk exposures is their changing portfolio weights.

Finally, we analyze the implications of our method for performance measurement. We find, similar to Ferson and Schadt (1996), that annualized alpha for funds with significant time-varying factor exposures rises on average by one percentage point when estimated using our model rather than the constant model. However, this finding masks much bigger changes at the individual fund level—we find a mean absolute difference of 2.7–4.6% between annualized alphas estimated using the constant model and our three time-varying exposure models.

The outline of the paper is as follows. Section I situates our paper in the literature on the dynamic performance evaluation of managed investments. Section II describes our modeling approach and Section III describes the data used in our analysis. Section IV presents analyses that verify our proposed method works well in practice, and Section V presents our main empirical results. Section VI looks at the sources of variation in hedge fund risk exposures and Section VII concludes.

I. Related Literature

Our paper contributes to the literature on dynamic performance measurement for actively managed investment vehicles, a topic that has recently experienced renewed interest. For example, Mamaysky, Spiegel, and Zhang (2008) use a Kalman filter–based model to track mutual fund risk exposures as latent random variables. Bollen and Whaley (2009) consider this approach for hedge funds, but instead recommend the use of optimal change-point regressions (as in Andrews, Lee, and Ploberger (1996)) to estimate structural breaks in hedge fund factor loadings. The change-point approach models risk exposures as constant between change-points, with abrupt changes to a new value at the change-points. The model pinpoints the time at which risk exposures change, although it is unable to provide insights into the underlying economic drivers of these changes. Our model provides a simple but economically interpretable alternative to this change-point approach in which time-varying betas are functions of observable conditioning variables.[4] The intellectual predecessor of our approach is Ferson and Schadt (1996), who use well-known predictors of returns as proxies for publicly available information and employ these instruments to estimate an unconditional version of their conditional model for the performance evaluation of mutual funds.[5]

[4] Patton and Ramadorai (2010) test the statistical performance of the approach in this paper as well as that of the change-point model on a large cross-section of hedge funds and funds of funds. They find that the "linear" model described in Section II.B.1 below yields better statistical performance than the change-point model, but that there are gains to combining the two approaches.

[5] Chen and Knez (1996) derive contemporaneous insights into conditional performance evaluation. These models are also related to Jagannathan and Wang (1996), who focus on risk adjustment

Our main contribution lies in the use of daily conditioning information to evaluate monthly reported performance. There have been other attempts to combine monthly returns and intra-monthly information to ascertain the higher-frequency variation in risk factor loadings, following an influential paper by Goetzmann, Ingersoll, and Ivkovic (2000), who show that Henriksson-Merton (1981) timing measures (discussed below) estimated from monthly data are biased in the presence of daily timing ability. Goetzmann, Ingersoll, and Ivkovic attempt to correct for this bias by cumulating daily put values on the S&P 500 for each month in their sample, which they incorporate as an additional regressor in their market-timing specifications (Ferson and Khang (2002) also present a conditional version of the holdings-based performance evaluation method that avoids the Goetzmann, Ingersoll, and Ivkovic bias). Ferson, Henry, and Kisgen (2006) consider an underlying continuous-time process for the term structure of interest rates to study monthly government bond fund performance, and find that time averages of daily interest rate movements are pivotal in explaining bond mutual fund performance. Although similar in spirit, our approach differs in a number of ways from the methods followed in these papers. First, our approach relies on intra-monthly products of factors and interaction variables, rather than on time-aggregated higher-frequency factors alone (which we additionally consider in the day-of-the-month model). Second, we posit several different daily models for hedge fund returns, which we use to uncover the actual intra-monthly *patterns* in hedge fund risk exposures. This allows us to economically interpret hedge funds' high-frequency risk exposure dynamics in addition to mitigating the inevitable loss of information that arises when attempting to infer intra-monthly dynamics using only monthly returns.

It is worth briefly mentioning a set of models that use conditioning information to detect time variation in managerial risk exposures in an attempt to find evidence of market-timing ability. One approach that is often employed (for example, by Treynor and Mazuy (1966) and Lehmann and Modest (1987) for mutual funds, and Chen and Liang (2007) for hedge funds) is to extend the standard single factor market model by including quadratic terms, or, as in Henriksson and Merton (1981), by interacting the market return with an indicator variable for the sign of the market return. Such regressions can be motivated using the model of Admati et al. (1986), in which a successful market-timing fund manager receives a noisy signal about the one-period-ahead market return—an idea that can be generalized to consider private signals about market attributes such as future market liquidity, as in Cao, Chen, and Liang (2009). As a consequence of the use of contemporaneous conditioning information, these models have two measures of managerial ability, namely, "timing," the coefficient on the interaction term between the factor and the contemporaneous variable representing the signal, and "selectivity," the intercept from the

for equities rather than performance evaluation. See also Ferson and Harvey (1991), and Evans (1994), among others. Mamaysky, Spiegel, and Zhang (2008) also find that adding observable variables to their model for mutual fund returns improves its performance, relative to a model in which only a latent factor drives variation in risk exposures.

unconditional estimation of the conditional model.[6] In contrast, in conditional performance evaluation models such as the one in this paper, the conditioning information is lagged, meaning that estimated alphas are measures of fund performance over and above those that can be obtained using public information signals and can be interpreted as measures of managerial ability in the usual manner.[7]

Finally, our use of daily returns on hedge fund indexes to validate our proposed method (see Section IV) adds to the sparse literature that uses daily data on investment managers' returns to measure their performance. Busse (1999) finds that mutual funds have significant volatility timing ability using daily returns data. Bollen and Busse (2001), also using daily data, confirm that mutual funds have significant market timing ability. Chance and Hemler (2001) use daily executed recommendations of market-timers and find that they have significant daily timing ability that vanishes when their performance is evaluated at the monthly frequency.

II. Modeling Time-Varying Hedge Fund Risk Exposures

In this section we first describe the conditional performance evaluation approach of Ferson and Schadt (1996), which is the initial point of departure for our model. We then present the three variants of our model that we estimate in Section V. To simplify the description of the various models we consider a simple one-factor model for capturing risk exposures, although in our empirical analysis in Section V we allow for multiple factors.

A. Models with Monthly Variation in Risk Exposures

Ferson and Schadt (1996) present a model in which betas evolve as a linear function of observable variables measured monthly:

$$r_{it} = \alpha_i + \beta_{it} f_t + \varepsilon_{it} \tag{1}$$

$$\text{where} \quad \beta_{it} = \beta_i + \gamma_i Z_{t-1}. \tag{2}$$

That is, the return on fund i is driven by a factor f_t, where the factor loading varies according to some zero-mean variable Z_{t-1}.[8] Substituting in the equation

[6] Holdings-based performance evaluation approaches have also been used to separate timing ability from selectivity (see Daniel et al. (1997), Chen, Jegadeesh, and Wermers (2000), and Da, Gao, and Jagannathan (2009)). Graham and Harvey (1996) use asset allocation recommendations in investment newsletters to evaluate whether they help investors time the market.

[7] Note that the approach in Ferson and Schadt (1996) is extended by Christophersen, Ferson, and Glassman (1998) to include the possibility of time-variation in alpha; this is also a possible extension to our approach.

[8] De-meaning Z_{t-1} ensures that we can interpret β_i as the average level of risk exposure. Using Z_{t-1} rather than Z_t means that we can interpret α_i as a measure of the fund's risk-adjusted performance, as per the discussion in Section I.

for β_{it} we obtain

$$r_{it} = \alpha_i + \beta_i f_t + \gamma_i f_t Z_{t-1} + \varepsilon_{it}, \tag{3}$$

which is easily estimated using OLS regression. Note that the constant-beta model is nested in the above specification, and the significance of time variation in beta for the i^{th} fund can be tested via a standard Wald test of the following hypothesis:

$$H_0^{(i)} : \gamma_i = 0 \ \text{ vs. } \ H_a^{(i)} : \gamma_i \neq 0. \tag{4}$$

B. Models with Daily Variation in Risk Exposures

Many hedge funds alter or turn over positions very frequently, thus it is possible that a hedge fund's risk exposure changes substantially within a month. This observation necessitates an extension of the above approach for modeling time-varying risk exposures. Consider the daily returns on hedge fund i, denoted r^*_{id}, and a corresponding daily factor model for these returns:

$$r^*_{id} = \alpha_i + \beta_{id} f^*_d + \varepsilon^*_{id}. \tag{5}$$

Let us assume that the factor loadings for this fund vary as a function of some conditioning variable Z, which is observable at a daily frequency:

$$\beta_{id} = g_i(Z). \tag{6}$$

We can consider various functional forms for $g_i(Z)$. The simplest is the direct analogue to the Ferson and Schadt (1996) approach, that is, where $g_i(Z)$ is linear in Z.

B.1. A Linear Model for Factor Exposures

To better understand how the linear model relates to Ferson and Schadt (1996), let Z^*_d denote the conditioning variable measured at the daily frequency and Z_d denote this variable measured at the monthly frequency (that is, Z_d will be constant within each month and jump to a new level at the start of each month). Then the linear model for $g_i(Z)$ can be written as

$$\beta_{id} = g_i(Z) = \beta_i + \gamma_i Z_{d-1} + \delta_i Z^*_{d-1}. \tag{7}$$

Substituting into (5) we obtain a simple interaction model for daily hedge fund returns:

$$r^*_{id} = \alpha_i + \beta_i f^*_d + \gamma_i f^*_d Z_{d-1} + \delta_i f^*_d Z^*_{d-1} + \varepsilon^*_{id.} \tag{8}$$

Returns on individual hedge funds are currently only available monthly, and thus to estimate this model we need to aggregate returns from the daily frequency up to the monthly frequency.[9] Define the monthly return on

[9] We use log returns, and so the monthly return is simply the sum of the daily returns. In this case, however, the linear factor model is only approximate. An alternative is to use simple

fund i as

$$r_{it} \equiv \sum_{d \in \mathcal{M}(t)} r_{id}^{*}, \tag{9}$$

where $\mathcal{M}(t)$ is the set of days in month t. Define f_t and Z_t similarly, and let n_t denote the number of days in month t. Then the specification for monthly hedge fund returns becomes

$$r_{it} = \alpha_i n_t + \beta_i f_t + \gamma_i f_t Z_{t-1} + \delta_i \sum_{d \in \mathcal{M}(t)} f_d^{*} Z_{d-1}^{*} + \varepsilon_{it}. \tag{10}$$

Note that the dependent variable above is now the monthly return on hedge fund i, and all variables on the right-hand side are also measured monthly. The new variable that appears in this specification relative to the Ferson-Schadt (1996) style specification in equation (3) is of the form $\sum f_d^{*} Z_{d-1}^{*}$. This is a monthly aggregate of a *daily* interaction term, and it captures variations in hedge fund risk exposures at the daily frequency. If the factor f_d^{*} and the conditioning variable Z_{d-1}^{*} are both available at the daily frequency, then under the assumption that ε_{id}^{*} is serially uncorrelated and uncorrelated with f_s^{*} for all (d, s) we are able to estimate the coefficients of this model using standard OLS. As above, for valid statistical inference we need to account for potential heteroskedasticity and nonnormality in the residuals. In Section I we present analyses based on real daily hedge fund index returns and simulated returns, both of which confirm that this modeling approach works well in realistic applications.

The constant-beta model is nested in the above specification, and the significance of time variation in beta can be tested via a standard Wald test of the following hypothesis:

$$H_0^{(i)} : \gamma_i = \delta_i = 0 \quad \text{vs.} \quad H_a^{(i)} : \gamma_i \neq 0 \cup \delta_i \neq 0. \tag{11}$$

Furthermore, we can test whether we find significant evidence of daily variation in hedge fund risk exposures, controlling for monthly variation, by testing whether the coefficient on the daily interaction term is zero:

$$H_0^{(i)} : \delta_i = 0 \quad \text{vs.} \quad H_a^{(i)} : \delta_i \neq 0. \tag{12}$$

Although it is expected that hedge funds adjust their risk exposures within the month, our ability to detect those changes depends on whether we can find observable daily interaction variables Z_d^{*} that are correlated with those changes. We pick four economically motivated Z variables in this paper, described in Section III.C below.

returns, making the factor model exact, but introducing an approximation error when aggregating to monthly returns. In the Internet Appendix we show that the approximation error introduced by both of these approaches is negligible for our data.

C. Day-of-the-Month Effects in Factor Exposures

Lakonishok et al. (1991) find that pension fund managers tend to increasingly sell losing stocks in the fourth quarter of the year, when funds' portfolios are closely examined by the sponsors. The authors suggest that this constitutes evidence of window-dressing, where managers alter their portfolios to impress sponsors.

Most hedge funds claim to generate "absolute returns," which are uncorrelated (i.e., zero beta) with widely used benchmarks. Akin to pension funds and mutual funds, hedge funds also periodically report their returns to publicly available databases.[10] Within these reporting intervals, hedge funds may pursue strategies that might not necessarily be zero- or even low-beta. Indeed, managers could pursue strategies that have low monthly-average betas on benchmarks despite having quite high betas within these months. One such strategy is for managers to have high exposures just following monthly reporting dates, and lower exposures just preceding the subsequent reporting date. Another possibility is that managers attempt absolute return strategies at the beginning of the month, and take leveraged positions on benchmarks at the end of the month so as to generate higher returns.

To detect such intra-month seasonalities, we consider the following specification:

$$\beta_{id} = g_i\,(Z) = \bar{\beta}_i \lambda(d, \theta_{\lambda,i}) + \gamma_i Z_{d-1} + \delta_i Z^*_{d-1}, \tag{13}$$

where the second and third terms in the expression are the same as in the linear model. For the leading term, $\lambda_i(d, \theta_\lambda)$, we follow Ghysels, Santa-Clara, and Valkanov (2006) and Andreou, Ghysels, and Kourtellos (2010) and use an "exponential Almon" function, which provides a flexible parametric function of the day of the month. We model this as a fraction of the month, d/n_t, to accommodate months with differing numbers of days:[11]

$$\lambda(d, \theta_{\lambda,i}) = \frac{\omega_{id}}{\sum_{d\in\mathcal{M}(t)} \omega_{id}}, \tag{14}$$

$$\omega_{id} \equiv \exp\left\{\theta_{i1}\frac{d}{n_t} + \theta_{i2}\left(\frac{d}{n_t}\right)^2\right\}. \tag{15}$$

Relative to the linear model, the $g_i(Z)$ function now includes an "intercept" that detects the variation in betas on specific days of the month. The Internet

[10] One difference, of course, is that hedge funds report only returns, not actual holdings. Another is that reporting for pension funds is mandatory, whereas it is discretionary for hedge funds. One potentially important implication of this discretionary reporting for our tests, highlighted by Getmansky, Lo, and Makarov (2004), is excessive hedge fund return smoothness. We control for this possibility using their suggested approach in our robustness checks.

[11] When the number of trading days per month is constant, the regression implied by equations (13) and (14) is a MIDAS regression, as in Ghysels, Santa-Clara, and Valkanov (2006) and Andreou, Ghysels, and Kourtellos (2010).

Appendix plots a range of possible shapes that can be taken by $\lambda\left(d, \theta_{i,\lambda}\right)$, which encompasses a number of economically interesting patterns including the possibility of window-dressing described above.[12] Importantly, this specification also includes the case in which $\lambda(d, \theta_{\lambda,i})$ is flat throughout the month, which indicates the absence of a day-of-the-month effect and thus allows us to formally test for the presence of these effects. In our empirical analysis we also consider the simple case of a *pure* day-of-the-month specification,

$$\beta_{id} = g_i\,(Z) = \bar{\beta}_i \lambda(d, \theta_{\lambda,i}), \tag{16}$$

and examine whether evidence exists for purely deterministic variation in risk exposures.

D. A Threshold Model for Factor Exposures

Both models presented above assume that beta variation is linear in the conditioning variable. We next consider a nonlinear description for these dynamics based on a threshold model, where beta is assumed to switch from one value to another once a threshold for the conditioning variable is reached:

$$\beta_{id} = \begin{cases} \beta_{i,lo}, & Z^*_{d-1} \leq \bar{Z} \\ \beta_{i,hi}, & Z^*_{d-1} > \bar{Z}. \end{cases} \tag{17}$$

For this model we use conditioning variables that cumulate gains, losses, or volatility through each month, and we choose the threshold $\bar{Z}$ as a function of the month-end values of the conditioning variable. For example, one conditioning variable that we consider here is the cumulated return on a market index. This specification posits that beta remains at one level so long as the cumulated market return remains above some threshold, but if the market return falls too far (relative to the average movement in market returns during a month) then the beta switches to a new level.[13] We also consider variables that cumulate volatility, or changes in liquidity, as described below.

III. Data

A. Hedge Fund and Fund of Funds Data

We use a large cross-section of hedge funds and funds-of-funds over the period from January 1994 to June 2009, which we consolidate from data in

[12] The Internet Appendix may be found in the online version of this article.

[13] We thank an anonymous referee for suggesting that we consider a threshold model. This specification is motivated by the idea that, when a hedge fund manager breaches a threshold for monthly profits/losses, the beta is set to a new level. As hedge fund returns are not available at the daily frequency, we cannot use them as daily conditioning variables. However, by using variables that correlate with hedge fund daily returns we hope to capture this effect if it is present in the data.

the HFR, CISDM, TASS, Morningstar, and BarclayHedge databases. The Internet Appendix contains details on the process followed to consolidate these data. The funds in the combined database come from a broad range of vendor-classified strategies, which are consolidate into 10 main strategy groups: Security Selection, Global Macro, Relative Value, Directional Traders, Funds of Funds, MultiProcess, Emerging Markets, Fixed Income, CTAs, and Other (which contains funds with missing vendor strategy classifications).[14]

Table I reports summary statistics on the hedge fund data. To overcome the well-known problem of return smoothing in monthly reported hedge fund returns, in our analysis we use "unsmoothed" returns which are estimated from raw returns fund-by-fund using the Getmansky, Lo, and Makarov (2004) moving average model with two lags (in the Internet Appendix, we verify our results using four lags as well as raw returns). Panel A of the table shows that the means of the reported returns and unsmoothed returns are similar, but as expected the distribution of the "unsmoothed" returns is slightly more disperse. The median fund has assets under management of USD 32MM, whereas the mean is much larger at USD 167MM, reflecting the significantly skewed size distribution that several other studies (Getmansky (2004) and Teo (2010)) have highlighted. The median management and incentive fees are 1.5% and 20%, respectively, consistent with earlier literature (Agarwal, Daniel, and Naik (2009)), and the withdrawal restrictions are also comparable to earlier literature (Aragon (2007)). Panel B of the table shows that the lengths of the return histories for the funds in the sample correspond closely to those reported by Bollen and Whaley (2009), with around half of our funds having five or more years of data available, and around 17% of our funds having less than three years of data. Finally, Panel C reports the distribution of funds across strategies: the two largest strategies are Security Selection (20.8%) and Funds of Funds (23.3%), whereas the two smallest strategies (not including Other, which captures those funds with unreported strategies) are Relative Value (1.0%) and Emerging Markets (3.4%). Given that our complete sample contains 14,194 individual funds, even the smallest strategy group has 146 distinct hedge funds, which enables us to undertake relatively precise strategy-level analyses.

B. Hedge Fund Factors

Throughout our analysis, the underlying factor model that we use is the seven-factor model of Fung and Hsieh (2004a,b). This model has been used in numerous previous studies (see Bollen and Whaley (2009), Teo (2009), and Ramadorai (2011)). The set of factors comprises the excess return on the S&P 500 index (SP500); a small minus big factor (SMB) constructed as the difference between the Wilshire small and large capitalization stock indexes; the excess

[14] The set contains both live and dead funds. The distribution of live versus defunct funds is roughly similar across the databases, and the total percentage of defunct funds is 46%, which is comparable to the 73% ratio reported in Agarwal, Daniel, and Naik (2009) although their sample period ends in 2002.

Table I

Summary Statistics

This table shows summary statistics for the funds in our sample. Panel A reports the percentiles of the pooled (cross-sectional) distribution of returns, unsmoothed returns, assets under management (AUM), management fees, incentive fees, and lockup and redemption notice periods. Panel B shows the percentages of funds in the consolidated sample of 14,194 funds that have return histories of the lengths specified in the column headers. Panel C shows the number and percentage of the 14,194 funds in each of the strategies represented in the rows.

Panel A

	Returns	Unsmoothed Returns	AUM ($MM)	Management Fee	Incentive Fee	Lockup (Days)	Redemption Notice (Days)
25th Prctile	−0.700	−0.841	9.400	1.000	10.000	0.000	10.000
50th Prctile	0.720	0.710	32.000	1.500	20.000	0.000	30.000
75th Prctile	2.230	2.363	106.756	2.000	20.000	90.000	45.000
Mean	0.845	0.847	166.714	1.484	15.162	94.176	33.913

Panel B

	<36 Months	≥36, <60	≥ 60
Length (return history)	17.423	31.062	51.515

Panel C

	Funds in Strategy	
	Percent	Number
Security selection	20.752	2,946
Global macro	6.243	846
Relative value	1.030	146
Directional traders	12.788	1,815
Funds of funds	23.341	3,313
Multi-process	12.520	1,777
Emerging markets	3.372	479
Fixed income	5.678	806
CTAs	13.973	1,983
Other	0.303	43
TOTAL	100	14,194

returns on portfolios of lookback straddle options on currencies (PTFSFX), commodities (PTFSCOM), and bonds (PTFSBD), which are constructed to replicate the maximum possible return to trend-following strategies on their respective underlying assets;[15] the yield spread of the U.S. 10-year Treasury bond over the three-month T-bill, adjusted for the duration of the 10-year bond (TCM10Y); and the change in the credit spread of Moody's BAA bond over the 10-year Treasury bond, also appropriately adjusted for duration (BAAMTSY).

[15] See Fung and Hsieh (2001) for a detailed description of the construction of these primitive trend-following (PTF) factors.

C. Variables Associated with Changes in Risk Exposures

We consider a small set of economically motivated variables to identify increases or decreases in hedge funds' exposure to systematic risks. These four variables correspond to four underlying drivers of managerial decisions to alter portfolio allocations, namely, liquidity, funding and leverage, volatility, and performance.[16]

There is growing recognition of the impact of liquidity on hedge fund and mutual fund performance. Pollet and Wilson (2008) document that mutual funds rarely diversify in response to increases in their asset base, and associate their result with limits to the scalability of fund portfolios, such as price impact or liquidity constraints. Aragon (2007) and Sadka (2009) find that liquidity risk is an important determinant of hedge fund returns, one that is not captured by the Fung–Hsieh (2004a) seven factors. Following Cao, Chen, and Liang (2009), we consider the case in which managers attempt to vary their exposure to risk factors in such a manner as to mitigate the influence of price impact—as liquidity rises (falls), we expect the absolute magnitude of risk exposures to rise (fall) as funds more (less) frequently enter or exit positions. To capture systematic time-series variation in liquidity at both the monthly and the daily frequencies, we employ the funding liquidity measure proposed by Garleanu and Pedersen (2011), namely, the TED spread (the three-month LIBOR rate minus the three-month T-bill rate).

Turning to the second possible driver, hedge fund managers' exposures to systematic risk factors will vary with the level of leverage they employ if their long and short positions do not exactly offset one another along the dimension of factor exposure (see Rubin et al. (1999), who document that hedge funds take on significant leverage). The leverage available to hedge funds will vary with the costs of borrowing, which we capture using the first difference of the level factor (the constant-maturity three-month U.S. T-bill rate).

The third variable that we consider is market volatility. As financial market volatility rises, hedge fund managers wishing to maintain fund return volatility constant may trim risk exposures, as in Ferson and Schadt (1996). We therefore include the VIX index (see Whaley (2000)), which is a measure of volatility extracted from the prices of options on the S&P 500 index. We also employ a measure of "realized volatility" (RV) based on intra-daily data on the S&P 500 index when we estimate the threshold model.[17]

Finally, hedge fund managers are often implicitly or explicitly benchmarked to commonly available indexes. When the returns on these benchmarks are high, managers may be tempted to increase their risk-factor loadings to avoid

[16] In an earlier version of this paper, Patton and Ramadorai (2010), we consider a set of 19 conditioning variables spanning the same four groups of variables, and we test the significance of the best-fitting conditioning variable for each fund using the "reality check" of White (2000).

[17] These realized volatilities are based on five-minute prices and obtained from the Oxford-Man Institute's "Realized Library" data, available at http://realized.oxford-man.ox.ac.uk/. See Heber et al. (2009) for details on how these measures are computed. The S&P 500 RV series is available from January 1996 until February 2009. Outside of this period, we use the simple squared returns as the volatility measure.

the perception that they are underperforming, and vice versa. With this in mind, we also include returns on the S&P 500 as a conditioning variable. All told, we have a set of four possible conditioning variables: dLevel, the TED spread, the S&P 500 return, and VIX. As the three variables other than the S&P 500 returns are highly serially correlated, we use a simple exponentially weighted moving average (EWMA) model (with the optimal EWMA smoothing coefficient estimated for each series using nonlinear least squares) to obtain their "surprise" component, and use this in place of the levels of these variables.[18]

IV. The Accuracy of Estimates of Daily Betas Using Monthly Returns

In this section we study the accuracy of our proposed method for estimating daily variations in the factor exposures of hedge funds using only monthly returns on these funds. We analyze this problem in two ways, and find support for our method in both cases. Although data on individual hedge fund returns are almost invariably available only at the monthly frequency, daily data on a collection of hedge fund style index returns have recently become available. These daily index returns are an ideal real-world data set on which to check the accuracy of our method. Our first approach is to employ the daily data on hedge fund index returns, and to compare the results obtained when estimating the model on daily data with those obtained when using our method on the monthly returns on these indexes. Our second approach, described in the Internet Appendix, is to conduct a simulation that is calibrated to match the key features of hedge fund returns, and study the performance of the proposed method under different features of the return-generating process.

A. Results Using Daily Hedge Fund Index Returns

Daily returns on hedge fund style indexes have recently become available from Hedge Fund Research (HFR); see Distaso, Fernandes, and Zikes (2009) for an analysis of these indexes. We use these data to check whether the estimates of hedge fund factor exposures that we obtain based only on monthly returns are similar to those that would obtain if daily data were available. As the HFR daily returns are available only at the index level and begin in April 2003, they are not a replacement for the comprehensive data that we employ on individual hedge funds. Nevertheless, this daily information provides us with valuable insights into the performance of our method.

We employ the daily HFR indexes for five hedge fund styles: equity hedge, macro, directional, merger arbitrage, and relative value.[19] The period

[18] The first-order autocorrelations of the resulting "surprise" components of these variables are 0.21, 0.20, −0.08, and 0.09 (monthly) and 0.12, −0.06, 0.02, and 0.10 (daily).

[19] In total there are nine HFR indices that are available for at least 24 months and that have a clear strategy definition. The four remaining style indices generate results that are similar to the included style indices; specifically, the convertible arbitrage and distressed securities indices have

April 2003 to June 2009 yields 1,575 daily observations and 76 monthly observations.[20] In our main empirical analysis in Section V below, we consider the seven-factor Fung–Hsieh model for hedge fund returns, but three of the Fung–Hsieh factors (the returns on three portfolios of lookback straddle options) are available only at the monthly frequency, and so are not suitable for our model of daily hedge fund index returns. We therefore restrict our attention to the four Fung–Hsieh factors that are available at the daily frequency. As in our main analysis below, we follow Bollen and Whaley (2009) and reduce the Fung–Hsieh model to a more parsimonious two-factor specification by using the Bayesian Information Criterion (equivalent in this application to maximizing the R^2 or adjusted R^2) to find the two Fung–Hsieh factors that best describe the daily hedge fund index returns. The chosen factors and the coefficients on these factors in constant parameter models using daily and monthly returns are presented in the Internet Appendix.

Table II presents the results of the linear model for time-varying factor exposures based on conditioning information, estimated using either daily returns or monthly returns. The models that are estimated are the two-factor versions of the models presented in equations (8) and (10):

$$\begin{aligned} r_{id}^* = \alpha_i + \beta_{i1} f_{1d}^* + \beta_{i2} f_{2d}^* + \gamma_{i1} f_{1d}^* Z_{d-1} + \gamma_{i2} f_{2d}^* Z_{d-1} \\ + \delta_{i1} f_{1d}^* Z_{d-1}^* + \delta_{i2} f_{2d}^* Z_{d-1}^* + \varepsilon_{id}^* \end{aligned} \tag{18}$$

$$\begin{aligned} r_{it} = \alpha_i n_t + \beta_{i1} f_{1t} + \beta_{i2} f_{2t} + \gamma_{i1} f_{1t} Z_{t-1} + \gamma_{i2} f_{2t} Z_{t-1} \\ + \delta_{i1} \sum_{d \in \mathcal{M}(t)} f_{1d}^* Z_{d-1}^* + \delta_{i2} \sum_{d \in \mathcal{M}(t)} f_{2d}^* Z_{d-1}^* + \varepsilon_{it}, \end{aligned} \tag{19}$$

where α_i is the daily alpha of the fund, β_{i1} and β_{i2} are the constant exposures to the factors f_1 and f_2, γ_{i1} and γ_{i2} capture variations in factor exposures that occur at the monthly frequency (with the variable Z_t), and δ_{i1} and δ_{i2} capture variations in factor exposures that occur at the daily frequency (with the variable Z_d^*).

If the methodology presented in Section I is accurate, then we would expect to see similar parameter estimates across the two sampling frequencies. Up to sampling variability, this is indeed what we observe: across all five indexes, the signs of the estimated coefficients generally agree, and cases of disagreement all coincide with at least one parameter estimate that is not significantly different from zero. As expected, the parameter estimates obtained from monthly returns are generally less accurate than those estimated using daily returns. Further supporting our approach, the p-values from the test for the significance of time-varying factor exposures agree in all but one case: for the equity hedge,

similar results to the relative value index, market neutral has similar results to macro, and event driven is similar to directional. These results are presented in the Internet Appendix.

[20] The HFR directional index started on July 1, 2004 and so slightly fewer observations are available for this series: 1,259 daily observations and 60 monthly observations.

directional, and relative value indexes, significant variation is detected using both daily and monthly returns; for the macro index, no significant variation is detected using either frequency; and for the merger arbitrage index, significant variation is found using daily data but not monthly data.

Table II also presents the correlation between the time series of daily factor exposures (betas) estimated using daily and monthly returns. For example, the correlation between the time series of daily exposure to the S&P 500 of the equity hedge index estimated using daily and monthly returns is 0.90, and the correlation of daily estimates of this index's exposure to BAAMTSY is 0.91. Similar positive results obtain for the directional and merger arbitrage indexes. For the macro index one of the correlations is negative while the other is positive, but for that index we obtain no evidence of time-varying beta using either daily or monthly returns (the p-values from the tests are 0.07 and 0.56, respectively), and so the estimated daily betas are essentially just noisy estimates of a constant value, in which case we would not necessarily expect a positive correlation between daily and monthly estimates. For the relative value index we find a positive correlation between the daily and monthly estimates of beta on the S&P 500 index, but a much lower correlation for the beta on the BAAMTSY index. The explanation for this lower correlation can be seen from the estimated values of γ_2 and δ_2: using daily data these are estimated to be positive and significant, whereas using monthly data they are negative and/or insignificant. In this case, the loss of precision from using monthly data may mean that in practice there are gains from setting insignificant parameters to zero.

In Figures 1 and 2 we illustrate the correspondence between the estimates of daily factor exposures estimated using actual daily index returns and monthly index returns, respectively. For clarity, we narrow the focus of these plots to the last quarter of our sample period (April 2009 to June 2009); similar conclusions can be drawn from other subperiods. The figures illustrate the strong similarity between the two estimates of time-varying exposure to the S&P 500 index, and provide further support for the modeling approach proposed in Section II.

We also conduct a simulation exercise, described in the Internet Appendix, that provides further evidence in support of the method. Overall, our analysis of daily returns on hedge fund indexes and our simulation results provide strong support for the reliability of our estimation procedure in practice. Given daily data on conditioning variables, the results of this section confirm that our method provides a means of obtaining reliable estimates of daily risk exposures from monthly hedge fund returns.

V. Empirical Evidence on Dynamic Risk Exposures

As described above, we follow Bollen and Whaley (2009) and reduce the full seven-factor Fung–Hsieh (2004a, 2004b) model, choosing a more parsimonious two-factor subset of these factors as the baseline static model. Figure 3 shows that the most frequently selected factor is the S&P 500 index, chosen for 60.5% of the funds. Of the remaining six factors, the most frequently selected is

Table II

The Linear Model for Daily and Monthly Hedge Fund Style Indexes

Table II shows results from a two-factor model applied to five hedge fund style index returns, identified in the first row of the table, allowing for time-variation in the factor exposures as a linear function of conditioning variables (i.e., $g(Z)$ is linear in Z, as in equation (7)). Two factors from the set of four daily Fung–Hsieh (2004a) factors are selected for each style using the Bayesian Information Criterion—these are identified in the Internet Appendix. The first row presents annualized alpha. Robust standard errors are reported below the parameter estimates, and the R^2 and adjusted R^2 are also reported. The p-values are reported for the joint significance of the coefficients on the interaction terms (Gamma1, Gamma2, Delta1, Delta2). The final two rows present the correlations between the time series of daily factor exposures estimated using daily and monthly data, for each of the two factors. For ease of comparison, Z is set to be dLevel for all indexes.

	Equity Hedge		Macro		Directional		Merger Arbitrage		Relative Value	
	Daily	Monthly	Daily	Monthly	Daily	Monthly	Daily	Monthly	Daily	Monthly
alpha	1.406	3.064	3.395	3.200	3.226	7.612	4.792	4.450	0.652	3.980
s.e.	2.217	2.053	3.782	3.859	3.529	3.213	1.417	1.402	2.860	1.820
beta1	0.308	0.327	0.066	0.121	0.318	0.248	0.092	0.104	0.107	0.076
s.e.	0.008	0.051	0.027	0.144	0.012	0.081	0.005	0.035	0.010	0.045
beta2	−1.873	−1.802	−0.748	−0.766	−2.483	−3.293	−0.677	−0.614	−1.900	−5.795
s.e.	0.291	0.698	0.244	1.248	0.470	1.037	0.186	0.477	0.375	0.619
gam1	0.006	0.010	0.001	0.003	0.005	0.002	−0.001	0.009	0.005	−0.010
s.e.	0.001	0.006	0.003	0.025	0.001	0.009	0.000	0.004	0.001	0.005
gam2	−0.009	0.027	0.039	0.069	0.004	−0.025	−0.059	0.027	0.129	0.002
s.e.	0.031	0.113	0.032	0.148	0.048	0.170	0.020	0.077	0.041	0.100
delta1	0.006	0.034	−0.012	−0.022	0.004	0.054	−0.003	−0.003	0.017	0.050
s.e.	0.002	0.012	0.008	0.071	0.003	0.018	0.002	0.008	0.003	0.011
delta2	0.221	0.325	−0.272	1.082	0.295	1.119	−0.006	0.187	0.100	0.099
s.e.	0.072	0.763	0.108	0.763	0.103	1.130	0.046	0.521	0.093	0.676
R2	0.575	0.734	0.019	0.049	0.470	0.715	0.291	0.262	0.131	0.846
R2adj	0.573	0.710	0.016	−0.035	0.468	0.683	0.288	0.197	0.127	0.832
pval	0.000	0.008	0.074	0.557	0.000	0.035	0.010	0.150	0.000	0.000
Corr-b1	0.897		0.985		0.420		−0.492		0.199	
Corr-b2	0.911		−0.776		0.992		−0.523		0.352	

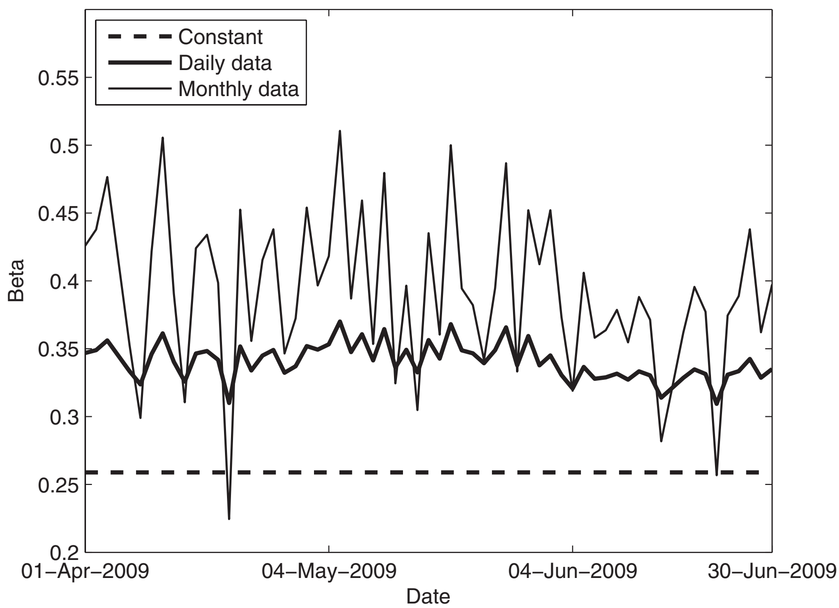

Figure 1. Estimated daily beta on the S&P 500 index for the Equity Hedge index. This figure depicts estimates of the daily exposure of the HFR equity hedge index to the SP500 index over the second quarter of 2009 from three models: constant beta, time-varying beta using daily returns on the index, and time-varying beta using the proposed method based only on monthly returns. The full-sample correlation between the estimates of time-varying beta using daily and monthly data is 0.90.

the size factor (SMB) while the second most frequently selected factor is the default spread (BAAMTSY), which are chosen for 33.1% and 32.6% of funds, respectively. The Internet Appendix breaks this down across the nine strategy groups and shows that the selected second factors are generally consistent with intuition about the factors on which different strategies load. With these "optimal" two-factor models for each individual fund, we now turn to the results from the different models for dynamic exposures to these factors.

A. Evidence of Time-Varying Risk Exposures

A.1. The Linear Model

Table III presents the results of statistical tests for our linear model of time-varying risk exposures, across the entire set of 14,194 funds in the database. For each of the four choices of conditioning variables we present the proportion of funds for which we can reject the null of constant factor exposures at the 5% level using different variations of the linear model.[21]

[21] All of the models we consider can be estimated using either ordinary or nonlinear least squares. However, our sample sizes are often short (we require only that a fund have at least

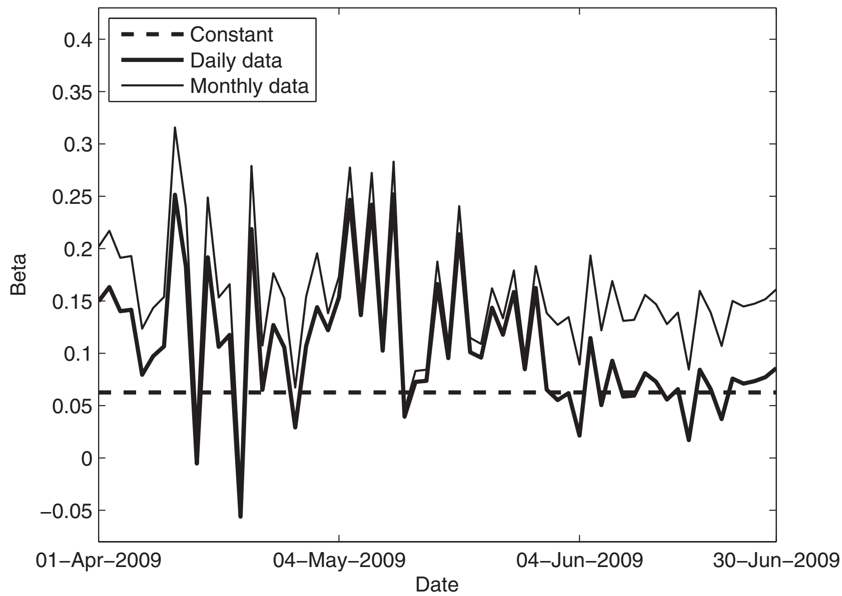

Figure 2. Estimated daily beta on the S&P 500 index for the Relative Value index. This figure depicts estimates of the daily exposure of the HFR relative value index to the S&P 500 index over the second quarter of 2009 from three models: constant beta, time-varying beta using daily returns on the index, and time-varying beta using the proposed method based only on monthly returns. The full-sample correlation between the estimates of time-varying beta using daily and monthly data is 0.20.

The first row of Table III presents the results from tests based on our proposed linear approach. The four different choices for conditioning variable that we employ yield a similar proportion of funds that can reject the null of constant factor exposures at the 5% level of significance. On average across choices of conditioning variable, this proportion is 22.4% of the total, or 3,180 funds. With a 5% level test we expect around 5% rejections even when the nulls are true, and it would be useful to know whether the proportion 22.4% is significantly greater than 5%. Establishing this requires an assumption on the correlation between the test statistics. If we assume that each fund's test statistic is independent of the others, then we can use the 95% upper quantile of the binomial distribution, and find that the critical value for the proportion of "significant" funds is 5.31%. A more conservative assumption on the correlation between test statistics of 20% leads (via simulation) to a critical value for the proportion of 11.1%, suggesting that our proportion of 22.4% is indeed significantly

24 observations to be included in our analysis) and as such heteroskedasticity- and autocorrelation-robust (HAC) standard errors such as Newey-West (1987) often do not perform well with such short samples; see Cochrane (2001, Chapter 11) for a discussion. We instead use a simple block bootstrap approach, based on Politis and Romano (1994) with an average block size of three, throughout our empirical analysis.

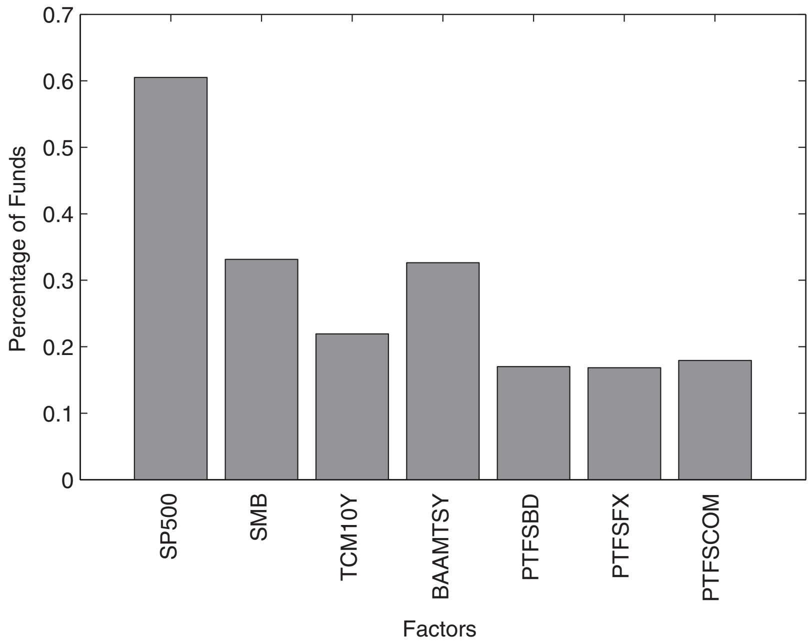

Figure 3. Fung–Hsieh Factor Selection. This figure depicts the proportion of times each of the Fung–Hsieh (2004) factors are selected in a two-factor model, as a percentage of the 14,194 individual funds.

higher than we would expect if none of the funds have significant time variation.[22]

In the remaining rows of Table III, we seek to isolate the sources of information from our modeling approach. In the second row, we test for time-varying risk exposures using only *monthly* conditioning information. That is, we force the coefficients on the daily information to be zero, and use a pure Ferson-Schadt (1996) approach. On average across conditioning variables, we can reject constant risk exposures for 12% of funds, which is substantially less than the 22.4% obtained when we combine daily and monthly information. The importance of daily information is reinforced by the next two rows of the table. When we test for the significance of daily information controlling for monthly information, we continue to find significance in approximately 22% of funds.

[22] We also undertook an alternative analysis of this issue using the false discovery rate (FDR) method of Barras, Scaillet, and Wermers (2010). The FDR approach provides an estimate of how many funds truly have time-varying risk exposures, but fail to reject the null in a statistical test of this hypothesis (due to a Type II error). For the results presented in this section we estimate that the true proportion of funds with time-varying risk exposures is 51.6%, averaging across the four conditioning variables (the individual proportions range from 48% to 53.0%). The FDR approach also provides a means of estimating the proportion of funds with significant results due purely to luck (Type I errors). In our application this is 2.4%, averaging across the four conditioning variables.

Table III

Hedge Fund Time-Varying Risk Exposures: Linear Model

This table shows results from the linear model (i.e., $g(Z)$ is linear in Z, as in equation (7)) for time-varying risk exposures in hedge fund returns. The table shows the proportion of 14,194 hedge funds for which we can reject the null hypothesis of constant risk exposures at the 5% level of significance, according to various models: the model with both daily and monthly Z variation allowed; the model with only monthly variation in Z; the incremental significance of daily Z variation over and above monthly Z variation; and the incremental significance of monthly Z variation over and above daily Z variation. The column headers list the four Z variables that we consider (the change in the three-month U.S. T-bill rate (labeled dLevel), S&P 500 returns, exponentially detrended VIX, and the TED spread), followed by a simple average of the proportions across Z variables.

	Conditioning Variables				
Model	dLevel	SP500	VIX	TED	Average
Daily and monthly	25.673	22.973	21.566	19.281	22.373
Monthly only	15.393	13.608	9.448	9.531	11.995
Daily given monthly	23.018	21.399	23.056	20.870	22.086
Monthly given daily	15.325	10.386	13.434	9.947	12.273

However, when we test for the significance of monthly information controlling for daily information, the proportion of significant funds drops to 12.3%. These results point to the importance of daily information for models of the dynamics of hedge fund risk exposures.[23]

Table IV Panels A and B break down the results from the linear model by the style of hedge fund. The table reveals interesting differences across conditioning variables. For example, the TED spread is the best conditioning variable for the Fixed Income style, whereas dLevel is the best for the Directional Traders and Security Selection styles, suggesting the importance of leverage to funds in these styles. Panel B of Table IV confirms that, for the majority of styles, daily conditioning information remains significant even when controlling for monthly conditioning information.

There are also interesting differences across styles. For example, the two styles with the highest proportions of funds rejecting the null of constant risk exposures are funds of funds and multiprocess funds, whereas CTA and Global Macro funds seem to have lower proportions of rejections. We attribute this difference to three main reasons. First, the returns of funds of funds and multiprocess funds are essentially an average across the returns of the multiple funds/strategies that they manage. This averaging makes their returns less noisy, allowing methods such as ours to more easily identify movements in risk exposures. (It is worth highlighting that such averaging is likely to be helpful

[23] In the Internet Appendix we present a series of robustness checks of these results. Specifically, we look at the sensitivity of our results to subsamples (1994 to 2001, 2002 to 2009); the number of lags used in the GLM "unsmoothing" model; the number of observations available on the fund; and the average size of the fund. Our results survive all of these checks. Our findings are strongest for the latter subsample, for funds with greater assets under management, and, not surprisingly, for funds with a longer history of available data.

Table IV

Linear Model Results by Hedge Fund Style

This table shows the proportion of hedge funds across the 10 styles for which we can reject the null hypothesis of constant risk exposures at the 5% level of significance, using the linear model (i.e., $g(Z)$ is linear in Z, as in equation (7)). Panel A shows these results for a model that allows for Z variation at both the daily and the monthly frequencies, and Panel B shows the proportion of funds for which allowing for daily variation is significant over and above a baseline monthly model. The columns show the four choices of Z considered and the average proportion across all four.

		Conditioning Variables				
Style	N(Funds)	dLevel	SP500	VIX	TED	Average
Panel A: Daily and Monthly						
Security selection	2,942	22.175	18.243	16.082	15.409	17.977
Macro	885	14.400	15.333	12.267	11.733	13.433
Relative value	146	19.118	19.118	20.588	18.382	19.302
Directional traders	1,813	24.563	22.479	16.958	16.789	20.197
Fund of funds	3,309	42.853	36.775	38.853	29.271	36.938
Multi-process	1,775	25.382	23.325	21.915	21.093	22.929
Emerging	478	13.763	19.140	18.065	12.043	15.753
Fixed income	805	16.255	21.456	13.914	23.667	18.823
CTA	1,981	12.171	9.013	9.737	10.132	10.263
Other	43	17.949	12.821	17.949	17.949	16.667
Panel B: Daily Given Monthly						
Security selection	2,942	18.101	15.551	17.641	16.082	16.844
Macro	885	11.200	9.467	11.867	11.200	10.933
Relative value	146	18.382	26.471	20.588	23.529	22.243
Directional traders	1,813	19.493	16.394	18.085	19.268	18.310
Fund of funds	3,309	40.217	35.876	41.178	33.488	37.690
Multi-process	1,775	24.442	24.031	24.266	23.090	23.957
Emerging	478	13.978	15.914	17.204	13.763	15.215
Fixed income	805	19.766	28.088	16.515	22.887	21.814
CTA	1,981	9.145	8.355	10.263	8.421	9.046
Other	43	17.949	12.821	12.821	12.821	14.103

in the case in which risk exposure variation is correlated across funds. The fact that both fund of funds and multiprocess funds exhibit high levels of risk exposure variation appears to provide evidence in support of this conjecture). Second, we employ four broad interaction variables, which capture systematic trends rather than fund-specific reasons for portfolio allocation shifts. This makes our variables better for capturing the risks of diversified portfolios of hedge funds such as funds of funds and multiprocess funds. Third, hedge fund factor models attempt to capture funds' dynamic trading strategies, but are not designed to capture variation in portfolio weights *across* strategies, that is movements from one strategy to another. It may be the use that our approach allows traditional factor models to better pick up these sorts of changes as well.

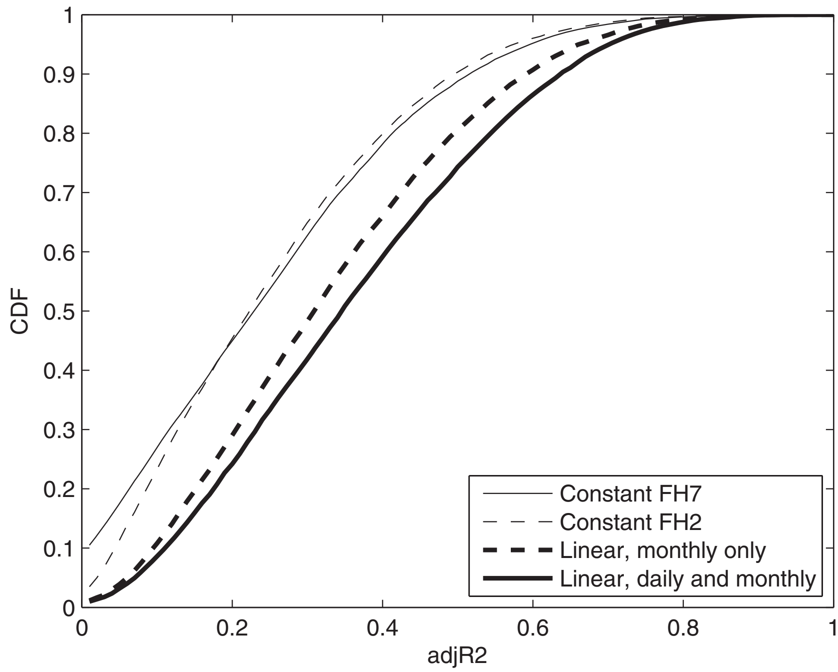

Figure 4. CDF of adjusted R^2 statistics from various models. This figure depicts the empirical cumulative distribution function (CDF) of adjusted R^2 statistics from the constant-parameter model using all seven factors, constant-parameter model using only two optimally chosen factors, the linear model with only monthly conditioning information, and the linear model with monthly and daily conditioning information and day-of-the-month effects, across 14,194 individual funds.

The importance of daily information for modeling hedge fund risk exposures is illustrated in Figure 4 for the entire set of 14,194 funds, plotting the cumulative distribution functions of the adjusted R^2 statistic for all funds for two static models (the best two Fung–Hsieh factors as well as the full seven-factor model) and two dynamic models (the linear model with monthly conditioning information only, and the linear model with both daily and monthly conditioning information). As Bollen and Whaley (2009) find for their change-point model, Figure 4 shows that our proposed models convincingly beat the constant-parameter models in the sense that the CDF of our models *everywhere* lie under those of the constant model. This is not only true for the model estimated using the two-variable subset of the seven Fung–Hsieh factors selected in our first-stage search procedure, but also for the full seven-factor model. The latter is an attempt to model dynamic risk exposures using an option-based replication of the returns from a posited dynamic trading strategy, and the graph confirms that this is insufficient on its own to capture movements in funds' risk exposures. Comparing the CDF for the adjusted R^2 statistics from the model using only monthly information to that from the model using both daily and monthly information, we find a substantial improvement: the CDF of the latter again

Table V

Mutual Fund Time-Varying Risk Exposures: Linear Model

This table shows results from the linear model (i.e., $g(Z)$ is linear in Z, as in equation (7)) for time-varying risk exposures in mutual fund returns. Panel A (B) of the table shows the proportion of 17,886 "equity" mutual funds in the CRSP database with >50% average allocation to common or preferred stocks (15,027 "bond" funds with ≤50% average allocation to common or preferred stocks) for which we can reject the null hypothesis of constant risk exposures at the 5% level of significance, according to various models: the model with both daily and monthly Z variation allowed; the model with only monthly variation in Z; the incremental significance of daily Z variation over and above monthly Z variation; and the incremental significance of monthly Z variation over and above daily Z variation. The column headers list the four Z variables that we consider (the change in the three-month U.S. T-bill rate (labeled dLevel), S&P 500 returns, exponentially detrended VIX, and the TED spread), followed by a simple average of the proportions across Z variables.

	Conditioning Variables				
Model	dLevel	SP500	VIX	TED	Average
Panel A: Equity Funds					
Daily and monthly	17.685	18.733	7.340	19.537	15.824
Monthly only	17.492	13.894	9.684	18.228	14.824
Daily given monthly	8.941	13.337	4.477	12.013	9.692
Monthly given daily	16.633	10.234	5.757	17.720	12.586
Panel B: Bond Funds					
Daily and monthly	13.966	15.756	6.759	13.033	12.379
Monthly only	12.078	14.899	5.386	12.429	11.198
Daily given monthly	10.422	9.719	5.667	10.539	9.087
Monthly given daily	10.808	12.049	4.110	14.286	10.313

lies everywhere below that of the former, confirming the importance of daily conditioning information for describing hedge fund risk exposures.

A.2. The Linear Model for Mutual Funds

The use of daily conditioning information is useful in performance evaluation for hedge funds, with their fast-moving positions and relatively unconstrained portfolios. Is daily conditioning information also helpful for mutual funds? Table V estimates the linear model using the returns of 32,913 mutual funds between 1994 and 2010 obtained from CRSP, which we divide into two groups, namely "equity funds" (those with over 50% average allocations to common or preferred stocks over their lifetimes) and the remainder, which we dub "bond funds."[24] For the baseline static factor model, we employ the usual Carhart (1997) four factors (the three Fama-French (1993) factors plus momentum),

[24] We select all funds from CRSP with at least 24 months of available returns data once they reach USD 15MM in total net assets and winsorize the monthly return data at the .01 and 99.99 percentile points of the pooled fund-month distribution to eliminate a few large outliers.

and augment these with TCM10Y and BAAMTSY to capture bond fund risk exposures.

Panels A and B of the table reveal interesting differences between conditioning variables. For instance, TED spread and dLevel are twice as useful as VIX for bond funds, and the S&P 500 returns are most informative for equity funds. However, the main insight from this table is that, for both groups of mutual funds, daily conditioning information, whereas helpful, is far less informative than for hedge funds. On average, the significance of daily conditioning information once we control for monthly information is between 9% and 10% for mutual funds, compared with 22% for hedge funds. These results are consistent with the quite different levels of flexibility built into the portfolios of hedge funds and mutual funds, and suggest that daily conditioning information is most important for managed portfolios such as hedge funds, which involve higher-frequency (as well as more flexible) trading strategies.

A.3. The Day-of-the-Month Model

We now consider a model that allows for *deterministic* variation in risk exposures, as a function of the day of the month. If hedge fund managers engage in window-dressing or if they aim to lock in a pre-specified return each month, then we may expect to see risk exposures varying systematically with the day of the month. Table VI presents results from estimating a model with *only* day-of-the-month effects, and compares the performance of this model to an augmented model that includes exposures that vary linearly with daily and monthly conditioning variables.[25]

For comparison, the first row of Panel A of the table reproduces the numbers from Panel A of Table III on the linear model. The second row of Panel A of Table VI shows that pure deterministic effects in hedge fund risk exposures are significant for 23% of funds, a relatively high proportion. The next row of the table shows that the day-of-the-month effects are even more significant when we control for daily and monthly conditioning information using the linear model described above. That is, by controlling for the stochastic variation in beta that arises through dependence on daily and monthly information, we are better able to detect the deterministic component of these variations.

Panel B of Table VI shows that the adjusted R^2 gain from the addition of day-of-the-month effects is substantial. Over the baseline linear model, the average improvement in adjusted R^2 is 12%. The second row of the same panel shows that a model with deterministic variation in risk exposures alone delivers adjusted R^2 statistics that are quite close to those of the linear model on average, suggesting the importance of these deterministic effects. Figure 5 again depicts the performance of various models graphically for the entire set of funds in the data, plotting the CDFs of the adjusted R^2 statistic for all funds

[25] Due to the increased complexity of this model, we increase the minimum number of observations required for inclusion in this analysis from 24 to 48, reducing the number of funds from 14,194 to 9,239.

Table VI

Hedge Fund Time-Varying Risk Exposures: Day-of-the-Month Model

This table shows results from the day-of-the-month model for time-varying risk exposures in hedge fund returns. For comparison, the top row of each panel reports results for the linear model (where $g(Z)$ is linear in Z, as in equation (7)) from Table III. The second row presents the "pure" day-of-the month model (where $g(Z)$ depends only upon the day of the month, as in equation (16)). The third row presents results from a combination of these models (where $g(Z)$ is a function of the day of the month and conditioning variables, as in equation (13)). Panel A of the table shows the proportion of individual hedge funds for which we can reject the null hypothesis of constant risk exposures at the 5% level of significance, according to the various models, with the third row reporting the incremental significance of the day-of-the-month variation over and above the linear model. The column headers list the four Z variables that we consider (the change in the three-month U.S. T-bill rate (labeled dLevel), S&P 500 returns, exponentially detrended VIX, and the TED spread), followed by a simple average of the proportions across Z variables. Panel B shows the mean adjusted R^2 statistics from the linear model; the model with only day-of-the-month effects; and the combined linear model with day-of-the-month effects, computed for all four choices of Z, followed by the average adjusted R^2 statistic across Zs.

	Conditioning Variables				
Model	dLevel	SP500	VIX	TED	Average
Panel A: Proportion of Funds with Significant Time-Varying Betas					
Linear	25.673	22.973	21.566	19.281	22.373
Day of the month effects only	23.421	23.421	23.421	23.421	23.421
Day of the month given linear	32.188	31.328	28.994	31.294	30.951
Panel B: Mean Adjusted R^2 Statistics					
Linear	28.940	28.680	28.420	28.050	28.523
Day of the month effects only	26.130	26.130	26.130	26.130	26.130
Linear with day of the month effects	32.548	32.029	31.617	31.662	31.964

from the different models that we consider. The graph shows that, whereas the day-of-the-month model has relatively high performance on average, the CDF of the adjusted R^2 statistics from this model everywhere lies above that of the linear model. That noted, the model with both day-of-the-month effects and linear dependence of risk exposures on conditioning variables dominates all of the other models.

Figure 6 shows the estimated shapes of the day-of-the-month effects from the different models. To summarize the dynamics of all of the funds in the data, we pick the median parameters across all funds with significant time-variation in risk exposures. The dashed line in the figure shows the shape from the model when only deterministic variation in beta is considered, whereas the solid line shows the shape when both deterministic and linear effects are allowed in the model. The figure shows that, in the pure deterministic version of the model, risk exposures are initially low, rising to a peak of 1.7 eight days following the end of the previous month, and declining rapidly from that point towards a level of 0.1 at the end of the month. When conditioning variables are included in the model, the pattern becomes less extreme but has essentially the same

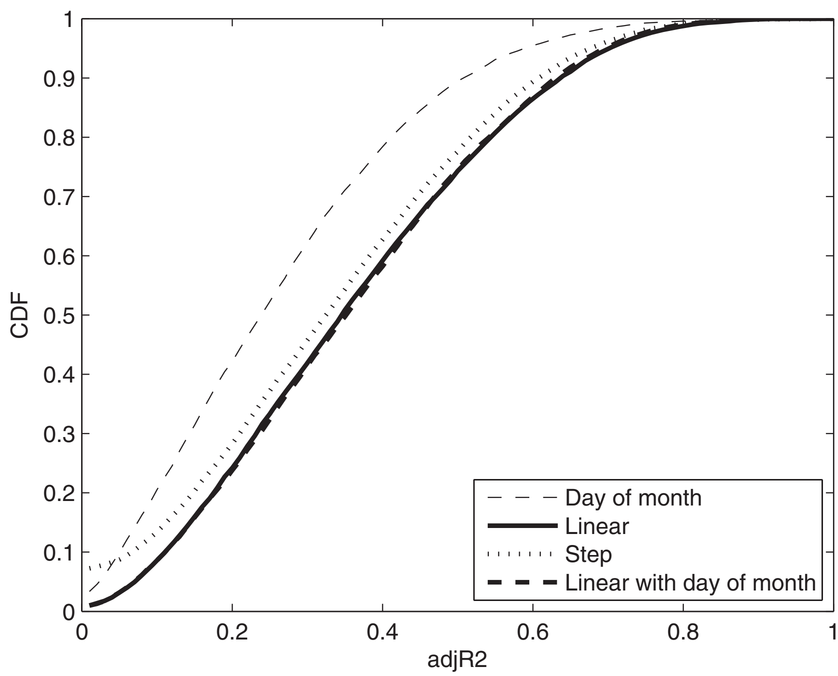

Figure 5. CDF of adjusted R^2 statistics from various models. This figure depicts the empirical cumulative distribution function (CDF) of adjusted R^2 statistics from the model with only day-of-the-month variation in risk exposures, the linear model with monthly and daily conditioning information, the threshold model, and the model with monthly and daily conditioning information as well as day-of-the-month effects, across 14,194 individual funds.

shape: high during the early part of the month, and declining significantly towards the end of the month.[26] Although this finding is striking, and suggests intra-month risk exposure alteration by hedge funds, other possibilities are also consistent with this pattern. For example, this pattern might be generated by intra-month deleveraging to satisfy end-of-month redemptions, although the fact that hedge fund assets under management (AUM) has been increasing over the sample period makes this an unlikely possibility (as end-of-prior-month inflows would be more consistent with rising risk exposures over the month). Another possibility is that the pattern may be generated by exposure to underlying assets, such as derivatives, that have risk exposures that vary predictably within a month. Even if this were the case, it would be puzzling because predictable exposures to such strategies are highly susceptible to front-running and "predatory trading" as in Brunnermeier and Pedersen (2005).

[26] A joint test that the average day-of-the-month patterns are the same across these two models yields a p-value of 0.08, suggesting no significant difference. This p-value is based on a test that takes the parameter estimates for each fund as given and is thus a lower bound on the true p-value: accounting for the estimation error in those estimates decreases the significance of this difference even further.

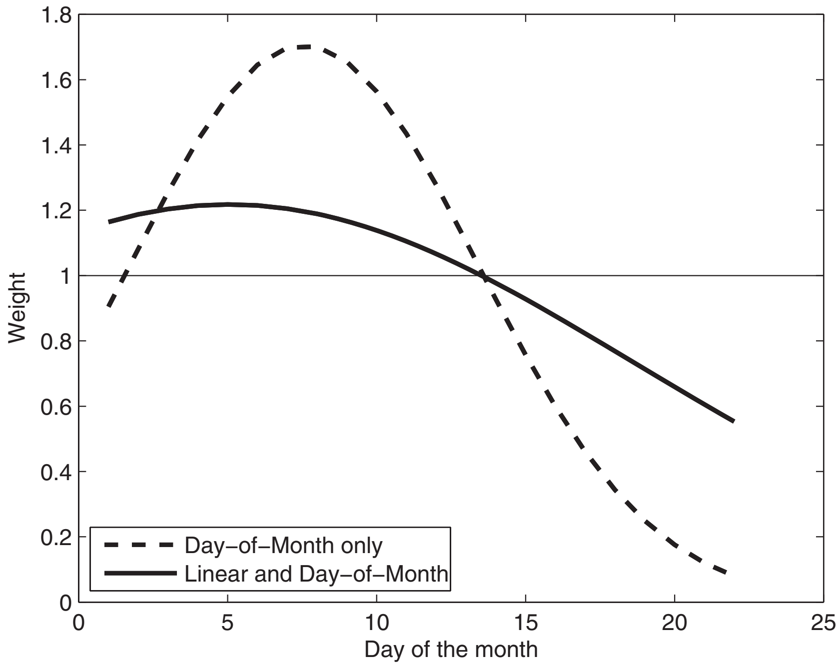

Figure 6. MIDAS weights on factor. This figure depicts day-of-the-month variation in risk exposures from the model with only day-of-the-month effects, and the model including both day-of-the-month effects and variation in risk exposures that is linear in both daily and monthly conditioning variables. The sample consists of funds with significant time-variation in risk exposures.

Table VII breaks the results of Table VI down by hedge fund style in an attempt to ascertain whether certain strategies are most often associated with such deterministic risk exposure profiles. Panel A of the table shows that, while it is indeed the case that high proportions of relative value funds such as option arbitrage and statistical arbitrage funds reject the null of constant risk exposures in favor of the pure day-of-the-month model, these funds by no means have a monopoly on deterministic variation in risk exposures. Panel B shows that, when linear variation in risk exposures is also permitted, every single one of the 10 styles has significant incremental time-variation explained by the addition of day-of-the-month effects. Indeed, this addition dramatically improves the proportion of Macro and CTA funds with significantly time-varying risk exposures, from 13.4% (10.3%) for Macro (CTA) funds using the baseline linear model to 24% (24%) when day-of-the-month effects are added to the baseline linear model. This suggests that the prevalence of important intra-month deterministic variation in risk exposures is not captured using the baseline linear model for these styles. This result is interesting, as it suggests either that there are intra-month deterministic movements in the risk exposures of the

Table VII

Day-of-the-Month Model Results by Hedge Fund Style

This table shows the proportion of individual hedge funds across 10 styles for which we reject the null of day-of-the-month variation in risk exposures at the 5% significance level. Panel A shows results for a model with day-of-the-month effects alone (as in equation (16)), and Panel B shows these results for the combined linear model with day-of-the-month effects (as in equation (13)). The columns show the four choices of Z considered and the average proportion across all four.

Panel A: Day-of-the-Month Effects Only

Style	N(funds)	Proportion Reject
Security selection	1,844	19.089
Macro	561	18.182
Relative value	87	34.483
Directional traders	1,187	19.377
Fund of funds	2,353	30.472
Multi-process	1,151	26.846
Emerging	293	23.891
Fixed income	533	26.266
CTA	1,212	17.327
Other	18	27.778

Panel B: Linear with Day-of-the-Month Effects

	Conditioning Variables				
Style	dLevel	SP500	VIX	TED	Average
Security selection	24.368	25.828	23.695	22.852	24.186
Macro	25.383	25.602	22.101	23.195	24.070
Relative value	31.325	31.325	32.530	34.940	32.530
Directional traders	23.484	23.313	21.435	22.289	22.630
Fund of funds	47.135	43.056	40.061	48.134	44.597
Multi-process	34.779	32.430	33.153	31.256	32.905
Emerging	28.322	32.517	21.678	25.874	27.098
Fixed income	29.377	32.296	25.486	27.043	28.551
CTA	24.236	23.443	23.556	24.575	23.953
Other	23.529	23.529	17.647	35.294	25.000

asset classes that Macro and CTA funds invest in, or that "exposure-dressing" is more prevalent for these funds.

A.4. The Threshold Model

The third model that we consider allows for *abrupt* movements in risk exposures when conditioning variables hit threshold values. To capture the idea that it is large market movements within the month that trigger such abrupt movements, we consider variables that accumulate through the month and reset at month-end. Three of the four variables used in the previous sections, dLevel, the S&P 500 index return, and the TED spread, all have natural

Table VIII

Insights from the Threshold Model

This table shows results from the threshold model (where $g(Z)$ is a nonlinear function of Z, as in equation (17), with risk exposures moving abruptly when Z hits a pre-specified threshold value) for time-varying risk exposures in hedge fund returns. The variables Z are named in the first column. The thresholds are quantiles, given in the second column, of the variable's month-end distribution. The third column of the table shows the proportion of 14,194 hedge funds for which we can reject the null hypothesis of constant risk exposures at the 5% level of significance, using this model. The fourth through seventh columns show the factor on which the loading changes according to the threshold model. The first row shows the unconditional average risk exposures to each of the factors. The next four rows show how these risk exposures change when each of the four conditioning variables crosses its threshold, falling below (rising above) the 10th (90th) percentile of its month end distribution at any point within the month. For example, the average fund's risk exposure to the S&P 500 falls by 64.4%, and its risk exposure to the credit spread rises (towards zero) by 107.5%, when S&P 500 realized volatility rises above the 90th percentile of its month-end distribution

		% Significant	SP500	SMB	TCM10Y	BAAMTSY
Average value			0.419	0.233	−1.780	−4.595
Cond. variable	Quantile		Percentage changes relative to average risk exposure			
dLevel	< 0.10	18.700	10.263	−19.742	88.202	45.680
S&P 500 return	< 0.10	15.499	−0.716	−126.180	−45.899	−35.930
S&P 500 volatility	> 0.90	35.000	−64.439	−16.738	72.753	107.552
TED	> 0.90	33.400	−52.745	−98.712	77.191	10.381

"cumulated" equivalents: we sum the daily values of these starting from the first day of the month, and reset them to zero at the start of the next month. The volatility variable used above, VIX, is not well suited to such cumulation, as it is a 22-day (option-implied) forecast of future volatility. Instead, we use S&P 500 index realized volatility and cumulate that over each month. In all cases we use as thresholds the 10% or 90% quantile of the distribution of month-end values for these series, which represent thresholds that signal particularly large movements in these series. For example, if the cumulated S&P 500 return falls below the 10^{th} percentile of monthly returns by, say, the 15^{th} of the month, then this model posits that the risk exposures of the fund jump to a new level. In "normal" months these thresholds will rarely be crossed, but when they are it is due to large movements in the conditioning variables.

Table VIII shows how hedge funds' risk exposures vary when the conditioning variables cross such thresholds. The first row of the table shows the average exposure to the factors listed in the columns, for example, the average static exposure to the S&P 500 across all funds is 0.419, and to the credit spread (BAAMTSY) is −4.595. The rows below show the exposure to these factors when the conditioning variable crosses the given threshold.

The first row of the table shows that, when dLevel hits the 10^{th} percentile of its month-end distribution, that is when the cost of leverage falls substantially, hedge funds increase their risk exposures to the S&P 500 (by 10%), reduce their exposures to small stocks (by 20%), reduce their short positions in credit spreads (a 46% increase towards zero) and long-term bonds, and increase their exposures to large stocks. The second row shows that, when the S&P 500 return hits low points within the month, hedge funds significantly move away from small stocks (turning their average positive exposure to SMB to a negative exposure). They also increase their negative exposures to long bonds and credit. When S&P 500 realized volatility spikes, hitting the 90^{th} percentile of its month-end distribution, hedge funds appear to cut their risk exposures across the board, moving significantly away from the S&P 500 and small stocks, and substantially decreasing their short positions in long bonds and especially credit spreads, trimming the latter almost to zero. Finally, when TED, our measure of funding liquidity, increases dramatically within the month (signifying a highly illiquid environment), hedge funds exit both S&P 500 stocks and small stocks, and again trim their short exposures to long bonds and credit.[27]

These results are interesting in light of the evidence provided at lower frequencies by Brunnermeier and Pedersen (2005), who document that hedge funds "rode" the technology bubble, and by and large fit the description of destabilizing rational speculation provided by DeLong et al. (1990), who show that directional trading by speculators in the presence of positive feedback traders can sometimes dominate trading strategies that take the opposite position to these noise traders. Furthermore, the evidence on the effects on dLevel and TED on hedge fund strategies adds to the somewhat sparse evidence on the effect of leverage and funding liquidity on hedge fund returns.[28] Such studies have been difficult to conduct given the lack of detailed data on this aspect of hedge funds' activities, and authors have adopted different strategies to capture these effects. For example, using simulations, Khandani and Lo (2011) highlight that systematic portfolio deleveraging by long-short equity hedge funds could have been responsible for the "quant meltdown" of August 2007.

B. Performance Measurement

Table IX shows how the use of our time-varying exposure model affects inferences about hedge fund alpha. The table contrasts the alphas obtained from a static factor model (using the optimally selected two factors from the set of all seven) with those obtained from the three time-varying exposure models that we consider. Across all funds, the average alphas from the three models look

[27] The indicators for the four threshold variables are correlated, as expected; however, only one pair is substantially correlated: the S&P 500 return/TED threshold variables have a correlation of 0.62; the next highest correlation (0.38) is between the dLevel and S&P 500 volatility thresholds. Note, however, that, even for the most highly correlated pair, the impacts of crossing each of these thresholds on risk exposures are different, suggesting that each of these threshold crossings carries some unique information.

[28] It is worth noting here that these two measures do not deliver identical results on hedge fund risk exposures. When the cost of leverage (i.e., dLevel) falls, hedge funds appear to trim small cap exposures marginally. However, it seems to be the case that shocks to funding liquidity (i.e., increases in TED) cause a significant retreat from small cap stocks.

Table IX

Accounting for Time-Variation in Factor Exposures: Implications for Alpha

This table presents analysis of the estimated alpha of individual hedge funds using a traditional static model (where risk exposures are assumed constant) presented in the first column, and three different models for time-varying risk exposures: Panel A presents results for the linear model (where $g(Z)$ is linear in Z, as in equation (7)); Panel B presents results for the linear model with day-of-the-month effects (where $g(Z)$ is linear in Z and also includes a deterministic day-of-the-month effect, as in equation (13)); Panel C presents results for the threshold model (where $g(Z)$ is a nonlinear function of Z, as in equation (17), with risk exposures moving abruptly when Z hits a pre-specified threshold value). The first two columns present annualized means of the estimated alphas from different models. The third column shows the cross-sectional mean difference between the dynamic model alpha and the static model alpha, the fourth column reports the cross-sectional mean of the absolute value of the difference of these alphas, and the final column reports the cross-sectional rank correlation of these alphas. Within each panel we present results separately for all funds, or only funds for which we can reject the null hypothesis of constant risk exposures. Note that the model in Panel B requires more time series observations, and so the universe of funds considered there is smaller than in Panels A and C (9,239 compared with 14,194 funds) and thus the results for the static model are different in that panel from the other two panels.

	Mean Alpha Static Model (1)	Mean Alpha T-V Beta Model (2)	Difference	|Difference|	Rank Correlation
		Panel A: Linear Model			
All funds	4.173	5.066	−0.893	3.170	0.868
t-stat	40.683	46.089	−16.733	67.635	60.964
Funds w/sig. variation	3.677	5.105	−1.428	3.497	0.806
t-stat	28.905	35.398	−17.478	49.191	55.928
		Panel B: Linear Model with Day-of-the-Month Effects			
All funds	4.331	5.035	−0.704	2.730	0.857
t-stat	44.662	48.644	−13.935	65.328	113.949
Funds w/sig. variation	3.818	4.852	−1.034	2.790	0.817
t-stat	32.667	38.270	−15.020	49.056	67.419
		Panel C: Threshold Model			
All funds	4.173	3.981	0.192	4.656	0.692
t-stat	40.683	35.452	2.292	62.729	40.036
Funds w/sig. variation	3.677	3.920	−0.243	4.189	0.674
t-stat	28.905	26.419	−2.027	38.957	20.602

quite similar, at around 4% per annum. Relative to the model with constant risk exposures, the linear model, with or without day-of-the-month effects, delivers alpha that is approximately 1% higher per annum than the constant model, whereas the threshold model delivers alpha that is marginally lower, by 20 basis points per annum.

When we restrict the sample of funds to those that reject the null of constant factor exposures, however, the alphas of these funds obtained from the time-varying exposure models are all significantly higher (by between 24 basis points and 1.4% per annum) than those from the static factor model. This suggests that, on average, hedge funds' variations in risk exposures are broadly beneficial to investors, a conclusion that is similar to that of Ferson and Schadt (1996) and to the more recent extension by Lo (2008), who shows that the "active component" of a manager's performance is attributable to the correlation between changing portfolio weights and asset returns.[29] Simply analyzing the average difference in the alphas misses an important point, namely, that the performance of some funds may improve and the performance of others may decline when the time-varying exposures model is applied. To account for this, that is to see if there are differences between the two models' inferences on any given fund, we measure the average of the absolute value of the difference between the alphas from the static and three dynamic models. The column labeled "| Difference |" shows that across all funds, we find a large and highly statistically significant difference of between 2.7% and 4.7% per annum (the latter being around the same size as the average static alpha) between the alphas from the two models. When we estimate this difference for only the funds that reject the null of constant factor exposures, we find that this difference remains roughly constant. Finally, we check the similarity of the performance rankings generated by the two models. The correlations of these rankings are high on average, at 81% across all funds and models. However, this masks variation across models, with rank correlations dropping to 67% depending on the model. This suggests that there are important differences in the relative performance evaluation of these hedge funds generated by the use of at least some variants of the time-varying risk exposure model.

VI. Decomposing Variation in Hedge Fund Risk Exposures

Our discussion in the previous sections explains patterns in hedge fund risk exposure dynamics as stemming from hedge fund manager portfolio allocation decisions. However, this is not the only source of possible variation in risk exposures. Changes in risk exposures could come from fund portfolio rebalancing, from variation in the underlying risk exposures of assets, or from the simultaneous occurrence of both. We, therefore, present a simple decomposition of the variation in hedge fund risk exposures into these separate components. Our decomposition is similar in spirit to Lo (2008), who decomposes variation in fund returns into variation in portfolio weights and variation in underlying asset returns, and shows that the correlation between changing portfolio weights and returns leads to time-varying exposures and excess expected returns. Our goal is somewhat different in this section, however, in that we focus on

[29] Titman and Tiu (2011) show that hedge funds that exhibit low (high) R^2 statistics on static factor models have higher (lower) Sharpe and information ratios. We find that funds with a high R^2 on the static model exhibit more time-variation in factors. This suggests that the measured Titman-Tiu performance differential might be smaller if time-variation in factor loadings were accounted for in alpha measurement. We check this, and confirm that the use of our time-varying factor model significantly reduces the Titman-Tiu performance differential for all funds in our sample, and virtually eliminates it for funds with significant time-variation in factor loadings.

understanding fund risk exposure variation rather than the underlying drivers of fund alpha.

For ease of exposition, consider the time-varying exposure of a fund to a single risk factor (β_t^f on factor f_t) as a function of the exposures of its holdings, labeled $\beta_{i,t}$. Writing $\omega_{i,t-1}$ for portfolio weights on assets $i = 1, \dots, n$ held by the fund, we obtain

$$\beta_t^f \equiv \sum_{i=1}^{n} \omega_{i,t-1}\beta_{i,t.} \tag{20}$$

We next rewrite the weights and stock betas in terms of deviations from their means, that is, $\omega_{i,t-1} = \bar{\omega}_i + \tilde{\omega}_{i,t}$ and $\beta_{it} = \bar{\beta}_i + \tilde{\beta}_{it}$. Then

$$\beta_t^f = \sum_{i=1}^{n} \bar{\beta}_i\bar{\omega}_i + \sum_{i=1}^{n} \bar{\beta}_i\tilde{\omega}_{i,t} + \sum_{i=1}^{n} \bar{\omega}_i\tilde{\beta}_{it} + \sum_{i=1}^{n} \tilde{\omega}_{i,t}\tilde{\beta}_{it}. \tag{21}$$

Decomposing variation in these risk exposures is then straightforward:

$$\begin{aligned}\mathrm{Var}\big[\beta_t^f\big] = {} & \underbrace{\mathrm{Var}\left[\sum_{i=1}^{n} \bar{\beta}_i\tilde{\omega}_{i,t}\right]}_{\text{Pure Weight}} + \underbrace{\mathrm{Var}\left[\sum_{i=1}^{n} \bar{\omega}_i\tilde{\beta}_{it}\right]}_{\text{Pure Beta}} + \underbrace{\mathrm{Var}\left[\sum_{i=1}^{n} \tilde{\omega}_{i,t}\tilde{\beta}_{it}\right]}_{\text{Weight-Beta}} \\ & + \text{covariance terms.}\end{aligned} \tag{22}$$

We label the different components by their source of variation, hence "pure weight," "pure beta," and "weight-beta," the last of which arises from the covariation between weights and betas. We treat the covariance terms as a residual and estimate its relative magnitude as well. We next use portfolio holdings information from 13-F filings to try to estimate the relative importance of these terms in driving changes in hedge fund risk exposures.

A. Hedge Fund 13-F Information

Data on hedge fund equity holdings are derived from the 13-F filings of hedge fund management companies. These 13-F filings only contain (large) long equity positions, and thus we focus on the 2,790 long/short equity funds in our sample of 14,194 funds. These 2,790 funds are associated with 1,754 unique management companies, which we match with all unique management companies filing 13-F reports. This procedure yields our final sample for this portion of the analysis: 252 matched long/short equity hedge fund management companies investing in 15,966 unique stocks between 1989Q1 and 2010Q4. For each of these stocks, we compute the time-varying quarterly beta using the four conditioning variables described above. To compute each management company's total beta, we apply the reported stock-level weights from the 13-F reports to the stock-level time-varying betas. Given this information, we are then able to attribute the variation in hedge fund risk exposures to their various sources.

Table X

Understanding the Sources of Time-Variation in Risk Exposures

This table decomposes the variation in betas for a set of 252 management companies. These companies include all the long/short equity hedge funds reporting quarterly 13-F filings information on their portfolio holdings that match our consolidated data set of hedge funds. The variation in total beta (estimated by applying reported stock-level weights from the 13-Fs to stock-level time-varying betas) is split into four components: pure weight variation (variation in portfolio weights holding underlying asset betas constant); pure beta variation (variation in underlying asset betas holding portfolio weights constant); weight-beta variation (arising from the co-variation between weights and betas); and a covariance term representing the difference between total beta variation and the sum of these three components, as in equation (22). Each column of the table reports the cross-sectional average variance of each of these components expressed as a fraction of the total variation in fund beta, with the associated White heteroskedasticity-robust *t*-statistic reported below each number. The rows represent the variable used to estimate time-varying underlying asset beta (at the quarterly frequency) in each case, namely the change in the three-month U.S. T-bill rate (labelled dLevel), S&P 500 returns, the first difference of VIX, and the quarterly change in the TED spread.

	Pure Weight	Pure Beta	Weight×Beta	Cov. Terms
Panel A: Full Sample				
dLevel	56.963	33.014	17.013	−6.991
t-stat	33.107	20.577	20.702	−3.352
SP500	79.606	9.848	18.560	−8.014
t-stat	38.087	8.696	17.493	−4.046
VIX	81.516	8.260	9.963	0.262
t-stat	58.984	8.878	13.448	0.239
dTED	72.902	15.321	18.221	−6.444
t-stat	41.099	7.677	13.863	−2.322
Average	72.747	16.611	15.939	−5.297
Panel B: Credit Crisis (2007–2010)				
dLevel	92.500	19.940	26.540	−38.980
t-stat	13.601	7.373	10.225	−4.584
SP500	82.506	18.330	23.383	−24.219
t-stat	28.064	8.101	8.854	−5.641
VIX	81.527	17.009	23.048	−21.584
t-stat	29.280	7.815	8.329	−5.322
dTED	79.946	17.681	25.357	−22.985
t-stat	30.457	7.202	8.289	−4.714
Average	84.120	18.240	24.582	−26.942

B. Results

Table X decomposes the variation in fund betas into its component parts. Panel A of the table implements the decomposition over the entire sample period, whereas Panel B implements it over the credit crisis (2007Q1 to 2010Q4) period.[30] Each column of both tables reports the cross-sectional average (across

[30] We compute this subperiod variance by taking deviations from the grand mean, to avoid introducing noise into the estimates stemming from estimation of means over short samples.

fund management companies) time-series variance of each of these components expressed as a fraction of the total variation in fund beta.

The table shows that the proximate source of variation in fund betas is portfolio weight variation, which accounts for 73% of total fund beta variation on average across conditioning variables over the entire sample, and a larger 84% share over the credit crisis period. Underlying asset beta variation accounts for a smaller fraction of the variation in total beta, a virtually invariant 16% to 18% depending on the period considered. The covariation between weights and betas jumps from 16% over the entire sample period to 25% of the total during the credit crisis, which might reflect attempts by hedge funds to anticipate movements in underlying asset risk exposures. The covariance terms, treated as a residual, account for a relatively low share of the variance over the full sample and a higher share during the crisis period.

Overall, the conclusion from this lower-frequency quarterly analysis of long equity holdings is that fund portfolio weight variation is the primary source of variation in observed fund risk exposures, accounting for roughly five times the variation generated by underlying asset betas. It should be noted here that the 13-F filings do not include hedge fund short positions; however, if anything this would bias the estimates of portfolio weight variation that we produce towards zero, because we are effectively seeing only half the variation in hedge fund portfolio weights by only looking at long positions.

VII. Conclusion

Recent research on hedge funds and mutual funds documents the importance of accounting for the dynamic nature of the risk exposures of these actively managed investment vehicles. Several approaches have been proposed in the literature, including modeling these risk exposures as unobserved latent factors and employing optimal change-point regression techniques. We add to this literature with a new conditional performance evaluation model that extends the approach pioneered by Ferson and Schadt (1996) to capture higher-frequency variation in hedge funds' factor exposures.

Using a comprehensive database of nearly 15,000 individual hedge funds over the 1994 to 2009 period, we find that our model performs well on statistical grounds, convincingly beating the constant parameter model as well as more sophisticated models. The extension of our model to capture daily variation in factor exposures is important in this context: the addition of daily information to a baseline model with purely monthly interaction variables allows us to identify 1,500 more funds with time-varying betas. This is in contrast with mutual funds: using data on nearly 33,000 individual mutual funds over the same sample period, we find that daily information adds relatively little beyond monthly information. This is consistent with hedge funds updating their positions at a higher frequency than mutual funds.

In addition to its good statistical performance, our approach is able to shed light on hedge fund behavior at horizons that have thus far proven resistant to scrutiny. We document significant intra-month variation in risk exposures, which are highest immediately following the end-of-month reporting date and then progressively lower, reaching a nadir just to the subsequent reporting date. We also find that hedge funds have a tendency to abruptly cut positions in response to significant market events such as sharp declines in market returns and liquidity, or abrupt increases in market volatility. These findings have interesting implications in light of recent public debates on the activities of hedge funds in financial markets.

Initial submission: April 2, 2010; Final version received: October 15, 2012
Editor: Jennifer Conrad

REFERENCES

Ackermann, Carl, Richard McEnally, and David Ravenscraft, 1999, The performance of hedge funds: Risk, return and incentives, *Journal of Finance* 54, 833–874.

Admati, Anat, Sudipto Bhattacharya, Paul Pfleiderer, and Stephen A. Ross, 1986, On timing and selectivity, *Journal of Finance* 41, 715–732.

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

Agarwal, Vikas, Naveen D. Daniel, and Narayan Y. Naik, 2009, Role of managerial incentives and discretion in hedge fund performance, *Journal of Finance* 64, 2221–2256.

Agarwal, Vikas, Naveen D. Daniel, and Narayan Y. Naik, 2011, Do hedge funds manage their reported returns? *Review of Financial Studies* 24, 3281–3320.

Agarwal, Vikas, William Fung, Yee Cheng Loon, and Narayan Y. Naik, 2011, Risk and return in convertible arbitrage: Evidence from the convertible bond market, *Journal of Empirical Finance* 18, 175–194.

Andreou, Elena, Eric Ghysels, and Andros Kourtellos, 2010, Regression models with mixed sampling frequencies, *Journal of Econometrics* 158, 246–261.

Andrews, Donald W. K., Inpyo Lee, and Werner Ploberger, 1996, Optimal changepoint tests for normal linear regression, *Journal of Econometrics* 70, 9–38.

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

Barras, Laurent, Olivier Scaillet, and Russ Wermers, 2010, False discoveries in mutual fund performance: Measuring luck in estimated alphas, *Journal of Finance* 65, 179–216.

Bollen, Nicolas P. B., and Jeffrey A. Busse, 2001, On the timing ability of mutual fund managers, *Journal of Finance* 56, 1075–1094.

Bollen, Nicolas P. B., and Veronika Krepely-Pool, 2009, Do hedge fund managers misreport returns? Evidence from the pooled distribution, *Journal of Finance* 64, 2257–2288.

Bollen, Nicolas P. B., and Robert E. Whaley, 2009, Hedge fund risk dynamics: Implications for performance appraisal, *Journal of Finance* 64, 985–1035.

Brunnermeier, Markus K., and Stefan Nagel, 2004, Hedge funds and the technology bubble, *Journal of Finance*, 59, 2013–2040.

Busse, Jeffrey A., 1999, Volatility timing in mutual funds: Evidence from daily returns, *Review of Financial Studies* 12, 1009–1041.

Cao, Charles, Yong Chen, and Bing Liang, 2009, Can hedge funds time market liquidity? Working paper, University of Massachusetts at Amherst.

Carhart, Mark M., 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57–82.

Chance, Don M., and Michael L. Hemler, 2001, The performance of professional market timers: Daily evidence from executed strategies, *Journal of Financial Economics* 62, 377–411.

Chen, Hsiu-Lang, Narasimhan Jegadeesh, and Russ Wermers, 2000, The value of active mutual fund management: An examination of the stockholdings and trades of fund managers, *Journal of Financial and Quantitative Analysis* 35, 343–368.

Chen, Yong, and Bing Liang, 2007, Do market timing hedge funds time the markets? *Journal of Financial and Quantitative Analysis* 42, 827–856.

Chen, Zhiwu, and Peter J. Knez, 1996, Portfolio performance measurement: Theory and applications, *Review of Financial Studies* 9, 511–555.

Christopherson, Jon A., Wayne E. Ferson, and Debra A. Glassman, 1998, Conditioning manager alpha on economic information: Another look at the persistence of performance, *Review of Financial Studies* 11, 111–142.

Cochrane, John. H., 2001, *Asset Pricing* (Princeton University Press, Princeton, NJ).

Da, Zhi, Pengjie Gao, and Ravi Jagannathan, 2009, When does a mutual fund's trade reveal its skill? Working paper, Northwestern University.

Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers, 1997, Measuring mutual fund performance with characteristic-based benchmarks, *Journal of Finance* 52, 1035–1058.

DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. 1990, Positive-feedback investment strategies and destabilizing rational speculation, *Journal of Finance* 45, 374–397.

Distaso, Walter, Marcelo Fernandes, and Filip Zikes, 2009, Tailing tail risk in the hedge fund industry, Working paper, Imperial College London.

Evans, Martin D. D., 1994, Expected returns, time-varying risk and risk premia, *Journal of Finance* 49, 655–680.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Ferson, Wayne E., and Campbell R. Harvey, 1991, Conditioning variables and the cross section of stock returns, *Journal of Finance* 54 1325–1360.

Ferson, Wayne E., Tyler Henry, and Darren Kisgen, 2006, Evaluating government bond fund performance with stochastic discount factors, *Review of Financial Studies* 19, 423–456.

Ferson, Wayne E., and Kenneth Khang, 2002, Conditional performance measurement using portfolio weights: Evidence for pension funds, *Journal of Financial Economics* 65, 249–282.

Ferson, Wayne E., Sergei Sarkissian, and Timothy Simin, 2008, Asset pricing models with conditional alphas and betas: The effects of data snooping and spurious regression, *Journal of Financial and Quantitative Analysis* 43, 331–354.

Ferson, Wayne E., and Rudi W. Schadt, 1996, Measuring fund strategy and performance in changing economic conditions, *Journal of Finance* 51, 425–461.

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

Fung, William, and David A. Hsieh, 2001, The risk in hedge fund strategies: Theory and evidence from trend followers, *Review of Financial Studies* 14, 313–341.

Fung, William, and David A. Hsieh, 2004a, Hedge fund benchmarks: A risk based approach, *Financial Analysts Journal* 60, 65–80.

Fung, William, and David A. Hsieh, 2004b, Asset-based style factors for hedge funds, *Financial Analysts Journal* 60, 65–80.

Fung, William, David A. Hsieh, Narayan Y. Naik, and Tarun Ramadorai, 2008, Hedge funds: Performance, risk and capital formation, *Journal of Finance* 63, 1777–1803.

Garleanu, Nicolae, and Lasse Heje Pedersen, 2011, Margin-based asset pricing and deviations from the law of one price, *Review of Financial Studies* 24, 1980–2022.

Getmansky, Mila, 2004, The life cycle of hedge funds: Fund flows, size and performance, Working paper, University of Massachusetts at Amherst.

Getmansky, Mila, Andrew W. Lo, and Igor Makarov, 2004, An econometric model of serial correlation and illiquidity in hedge fund returns, *Journal of Financial Economics* 74, 529–610.

Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov, 2006, Predicting volatility: How to get the most out of returns data sampled at different frequencies, *Journal of Econometrics* 131, 59–95.

Goetzmann, William N., Jonathan E. Ingersoll, and Zoran Ivkovic, 2000, Monthly measurement of daily timers, *Journal of Financial and Quantitative Analysis* 35, 257–290.

Graham, John R., and Campbell R. Harvey, 1996, Market timing ability and volatility implied in investment newsletters' asset allocation recommendations, *Journal of Financial Economics* 42, 397–422.

Hasbrouck, Joel, 2009, Trading costs and returns for U.S. equities: Estimating effective costs from daily data, *Journal of Finance* 64, 1445–1477.

Heber, Gerd, Asger Lunde, Neil Shephard, and Kevin Sheppard, 2009, Oxford-Man Institute's Realized Library, Version 0.1, Oxford-Man Institute, University of Oxford.
Henriksson, Roy D., and Robert C. Merton, 1981, On market timing and investment performance, II. Statistical procedures for evaluating forecasting skills, *Journal of Business* 54, 513–533.
Jagannathan, Ravi, Alexey Malakhov, and Dmitry Novikov, 2010, Do hot hands exist among hedge fund managers? An empirical evaluation, *Journal of Finance* 65, 217–255.
Jagannathan, Ravi, and Zhenyu Wang, 1996, The conditional CAPM and the cross-section of expected returns, *Journal of Finance* 51, 3–54.
Khandani, Amir E., and Andrew W. Lo, 2011, What happened to the quants in August 2007? Evidence from factors and transactions data, *Journal of Financial Markets* 14, 1–46.
Kosowski, Robert, Narayan Y. Naik, and Melvyn Teo, 2007, Do hedge funds deliver alpha? A Bayesian and bootstrap analysis, *Journal of Financial Economics* 84, 229–264.
Lakonishok, Josef, Andrei Shleifer, Richard Thaler, and Robert Vishny, 1991, Window dressing by pension fund managers, *American Economic Review Papers and Proceedings* 81, 227–231.
Lehmann, Bruce N., and David M. Modest, 1987, Mutual fund performance evaluation: A comparison of benchmarks and benchmark comparisons. *Journal of Finance* 42, 233–265.
Liang, Bing, 1999, On the performance of hedge funds, *Financial Analysts Journal* 55, 72–85.
Lo, Andrew, 2008, Where do alphas come from? A new measure of the value of active investment management, *Journal of Investment Management* 6, 1–29.
Mamaysky, Harry, Matthew Spiegel, and Hong Zhang, 2008, Estimating the dynamics of mutual fund alphas and betas, *Review of Financial Studies* 21, 233–264.
Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semidefinite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.
Patton, Andrew J., 2009, Are "market neutral" hedge funds really market neutral? *Review of Financial Studies* 22, 2495–2530.
Patton, Andrew J., and Tarun Ramadorai, 2010, On the dynamics of hedge fund risk exposures, Working paper, Oxford-Man Institute of Quantitative Finance.
Politis, Dimitris N., and Joseph P. Romano, 1994, The stationary bootstrap, *Journal of the American Statistical Association* 89, 1303–1313.
Pollet, Joshua, and Mungo I. Wilson, 2008, How does size affect mutual fund behavior? *Journal of Finance* 63, 2941–2969.
Ramadorai, Tarun, 2011, The secondary market for hedge funds and the closed-hedge fund premium, *Journal of Finance*, 67, 479–512.
Rubin, Robert E., Alan Greenspan, Arthur Levitt, and Brooksley Born, 1999, Hedge funds, leverage, and the lessons of long-term capital management, Report of the President's Working Group on Financial Markets.
Sadka, Ronnie, 2009, Liquidity risk and the cross-section of expected hedge fund returns, Working paper, Boston College.
Teo, Melvyn, 2009, Geography of hedge funds, *Review of Financial Studies* 22, 3531–3561.
Teo, Melvyn, 2010, Hedge fund capacity constraints, liquidity, and hierarchy costs, Working paper, Singapore Management University.
Titman, Sheridan, and Cristian Tiu, 2011, Do the best hedge funds hedge? *Review of Financial Studies* 24, 123–168.
Treynor, Jack L., and Kay K. Mazuy, 1966, Can mutual funds outguess the market? *Harvard Business Review* 44, 131–136.
Whaley, Robert E., 2000, The investor fear gauge, *Journal of Portfolio Management* 26, 12–17.
White, Halbert, 2000, A reality check for data snooping, *Econometrica* 68, 1097–1126.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's web site:

Appendix S1: Internet Appendix