Goldman Sachs Asset Management

Quantitative Investment Strategies

2018

An Alternative to Hedge Fund Investing
A Risk-Based Approach

1 Introduction

By virtue of their lack of investment constraints relative to traditional equity and fixed income managers, hedge funds have produced positive, diversifying returns for more than 20 years.[1] Investors have therefore used hedge funds to complement core equity and fixed income allocations with the expectation that this will result in an increase in overall portfolio efficiency. However, investing in hedge funds presents a distinct set of challenges for investors, notably liquidity restrictions, potential lack of transparency into the investment strategy, extensive due-diligence requirements as well as their fee structures. In the context of public equity market mutual funds, one response to some of those challenges has been to passively track a representative market benchmark. Unfortunately, the concept of the market portfolio as a representative benchmark, founded in the Capital Asset Pricing Model (CAPM) and Efficient Market Hypothesis (EMH), does not exist with hedge funds. Against this background, this article discusses an alternative to hedge fund investing. Informed by techniques from other asset classes, it outlines a factor-based approach to identifying the systematic risk exposures taken by hedge funds. These economically intuitive factors based on academic research are well-defined, liquid and can be implemented at relatively low cost. A portfolio of these systematic factors can provide investors with access to a hedge-fund-like return profile.

There are several reasons why a representative market benchmark does not exist for hedge funds. Leaving aside the fact that there is hardly a consensus definition of what a hedge fund is, it is impossible to passively track a benchmark representative of the entire hedge fund universe because of, among other issues, coverage restrictions of hedge fund data sources and investment frictions. On the one hand, hedge funds may report information to one or more of multiple hedge fund databases at their sole discretion, with the result that each database, and all databases collectively, provides only a partial representation of the hedge fund universe. On the other hand, the investment frictions associated with hedge funds (e.g. lockups, minimum investment amounts) and extensive due-diligence requirements represent significant barriers to initiate and maintain coverage of any sizeable and diverse portfolio of hedge funds, therefore posing further challenges to a passive investment approach.

[1] Hedge fund returns - as measured by the HFRI Fund Weighted Composite Index – returned a Sharpe ratio of 0.61 and an information ratio versus equities of 0.51 from September 1997 until September 2017, illustrating their ability to deliver strong returns in excess of the equity risk premium.

Given the lack of a viable hedge fund benchmark for investors to track passively, the question arises – is there a case to be made for a select portfolio of hedge funds instead. Investors may naturally strive to select those hedge funds which consistently and persistently produce diversifying and positive returns. In practice, the lack of transparency not only in the investment strategy but also in the reporting of hedge fund performance, positions, and attribution (which is often voluntary with no clearly defined standards in existence), can make it difficult to distinguish luck from skill. Additionally, this article quantifies the lack of performance persistence among hedge funds on a year-on-year basis. As outlined in Section 2.2, out of the top 20% funds in terms of past-year performance, only 29% of funds are found to be able to repeat this placement in the next year. This is in line with the academic literature on hedge fund manager performance persistence, as summarized for example by Agarwal et al (2015) and Eling (2009). While there may be a degree of persistence over a shorter-term horizon, i.e. periods of six months or less, this literature finds that the evidence for persistence becomes much more challenged over intermediate- to long-term horizons. This in turn implies that even if a hedge fund investor can continuously identify successful individual hedge funds ex ante they would be required to turn over their portfolio quite frequently. Additionally, subscription/redemption cycles as well as manager relationship constraints present material implementation challenges, leaving only potentially the most sophisticated investors with sufficient expertise and resources to dynamically adjust these types of portfolios.

The alternative investment approach proposed in this article acknowledges both the lack of a representative market benchmark as well as the challenges around maintaining a well-performing select portfolio of hedge funds. In order for investors to manage the dispersion in the performance of individual hedge funds, it argues in favor of a sufficiently diversified universe of hedge funds. While individual hedge funds may be highly idiosyncratic in their investment styles and resulting return profiles, such broadly diversified portfolios of hedge funds exhibit a higher degree of stability when it comes to the drivers of their return evolution over time. The discussed portfolio construction approach argues in favor of inferring such return drivers using systematic factor exposures of hedge funds, instead of the creation of large portfolios of direct hedge fund holdings. This study is grounded in the work of Fama and French (1992) on cross-sectional equity pricing and of Sharpe (1992) on asset-class factor models, Fung and Hsieh (1997, 2001, 2004) and Agarwal and Naik (2000a, 2000b, 2004), among others, that have pioneered this type of analysis of systematic return drivers for hedge funds.

The well-defined, liquid and relatively low cost factor exposures we employ fall into two categories, traditional and alternative risk premia. Traditional risk premia are individual “long only” market factors (betas) such as equities or fixed income. Alternative risk premia are defined as collections of investment rules and strategies that are often employed by hedge funds that can be implemented using liquid financial instruments and therefore have similar liquidity as traditional market factors. Particularly through its emphasis on alternative risk premia, the suggested methodology accomplishes enhanced tracking of the performance of a broad portfolio of hedge funds in comparison to, for example, Hasanhodzic and Lo (2007) or Hill et al (2004). Liquid access to these two categories of risk premia in an investment vehicle provides advantages over portfolios of individual hedge funds, and potentially even over investments in fund-of-hedge funds, such as liquidity, affordability, transparency and clear return attribution. A portfolio of these two categories of risk premia could be the solution for investors concerned about the challenge of performance consistency of portions of their hedge fund universe. Another advantage of such an investment approach is that it leaves open the possibility of investors to complement their portfolios with investments in specifically selected individual high-conviction hedge funds.

Individual high-conviction hedge funds might indeed be delivering attractive returns over and above the performance of traditional or alternative risk premia. This raises an important caveat about the investment approach to make the traditional and alternative risk premia exposures of hedge funds available to investors, as it does not provide access to the “unexplained” portion that may be present in the hedge fund universe beyond these systematic factor exposures. However, as outlined in Section 4.1, only 16% of the return of the hedge fund universe constructed for our analysis can be attributed to this unexplained portion. In turn, 84% of the return of the universe can be provided to investors by means of traditional and alternative risk premia. This percentage is not only due to static exposures to these risk premia but also captures time variation of hedge fund exposures to such risk premia, as the discussed methodology updates at regular intervals. Overall, the high degree of hedge fund performance capture translates into a correlation of 93.5% to the return time series of the underlying hedge fund universe.

While the proposed investment approach might represent a remedy for investors to the non-investability of a hedge fund benchmark, it is important to note that it behaves very differently from a passively tracking benchmark portfolio in the realm of, for example, public equity markets. Notwithstanding the very non-passive nature of the risk premia, particularly the alternative risk premia, the difference between the well-defined and liquid nature of the factors and the opaqueness and illiquidity of some hedge fund investment strategies will necessarily lead to a degree of tracking error. In the specific case of the proposed alternative to hedge fund investing, the tracking error amounts to approximately one third of the volatility of the hedge fund universe benchmark, per backtested analysis.

This document is structured as follows: Section 2 presents a more detailed introduction into the universe on which we base this analysis. We then analyze hedge fund performance persistence and elaborate on the similarities of portfolios of hedge funds of different sizes compared to the overall hedge fund universe. In Section 3, we present the set of traditional and alternative risk premia that allow us to identify the systematic drivers of hedge fund performance and discuss the weight estimation framework to allocate to those premia to emulate the risk-return characteristics of hedge funds in liquid form. Section 4 discusses the efficacy of the discussed weight estimation procedure. It further presents an explicit return and risk decomposition of overall hedge fund returns into traditional risk premia, alternative risk premia as well as an unexplained component. Section 5 complements the analysis with a cross-sectional analysis of the evolution of fees and liquidity of hedge funds. It also presents an outlook on the role that liquid tracking might be able to play against the background of recent developments in the hedge fund universe. Finally, Section 6 concludes with a perspective on the broader universe of liquid alternative investment vehicles that has emerged in recent years.

2 Benefits of a Diversified Portfolio of Hedge Funds

In this section, we focus on the hedge fund dataset that is at the core of the subsequent analysis of systematic performance drivers. We first describe the construction of the proprietary aggregate hedge fund data set and review its current and historic properties such as number of funds and assets under management.

We then focus on an analysis of performance persistence and highlight the lack thereof on a year-on-year basis. This lack of persistence suggests that hedge fund investors aiming at selecting top performing hedge funds would have to rebalance hedge fund portfolios more frequently and to a larger extent than is practically feasible, a concern that hedge fund investors may address by increasing their hedge fund portfolio's diversification.

However, another finding in this section is the degree of convergence between hedge fund portfolios and the overall studied hedge fund universe, even for hedge fund portfolios with a limited number of individual funds. Paired with the persistence result, this finding is the fundamental justification for the use of a broad and diverse set of hedge funds and their corresponding investment strategies to draw inferences about systematic hedge fund return drivers.

2.1 Hedge Fund Universe

We source hedge fund information directly from hedge fund database providers. Hedge funds or their management companies[2] typically provide information about hedge-fund-level monthly returns as well as assets under management (AUM) on a monthly basis, paired with a host of more qualitative information such as classification or their fee structure.

The universe of hedge funds this study is based on is constructed from data provided by two hedge fund database providers, Hedge Fund Research, Inc. and BarclayHedge, LLC. As of December 2017, these two databases provide us with access to close to 14,000 hedge fund time series.[3] As found in Joenvaara et al (2016), these two databases exhibit a high degree of complementarity. In order to ensure comparability of the hedge fund return time series, we restrict attention to US Dollar-denominated return time series and require all return information to be reported net of all fees. Using a proprietary merging algorithm,[4] we then construct a point-in-time representation of the hedge fund universe from the filtered raw information available from the hedge fund data providers.[5] Using this merging algorithm allows the analysis to be driven by a more comprehensive universe of hedge fund strategies while reducing noise in the analysis due to double-counting entries which appear multiple times across both databases.

Exhibit 1 shows the number of funds as well as the total AUM of this universe. For the past 10 years, its coverage in terms of the number of funds has remained fairly steady at around 3,500 funds, which, according to the Hedge Fund Research, Inc. (HFR) Global Hedge Fund Industry Report from the third quarter of 2017, represents slightly less than half of the number of funds commonly considered to be in the hedge fund universe. In terms of AUM, the universe has declined in the aftermath of the Global Financial Crisis, but has been steadily increasing since then. It now stands at around $1.7trn, which, as with the number of funds, is approximately half of the overall AUM managed within the hedge fund industry.[6]

[2] Disclosure to hedge fund databases is voluntary and one might express concerns about the selection bias inherent in hedge funds or their management companies deciding to be included in a hedge fund database or not. Reasons for reporting hedge fund returns to a database are manifold and include, amongst others, increased publicity, requests by investors or a perceived higher institutional quality. Implementing the methodology on as broad a hedge fund universe as possible makes the results robust to individual hedge fund managers stopping to report their returns and abates some concerns about the data's comprehensiveness.
[3] Note that the approximately 14,000 time series include the overlap of funds reporting to both databases as well as, for example, multiple share classes being reported for individual funds.
[4] The algorithm groups time series that exhibit a high degree of commonality to limit duplication. This way, we ensure that specific hedge fund's returns are not disproportionately represented in the universe by virtue of their reporting style or their reporting to both databases simultaneously.
[5] As we have access to point-in-time files from the database providers, we can rely on their information about hedge funds as being available at historic points in time to construct our aggregate database, which addresses concerns about survivorship biases. Prior to 2009, we rely on so-called graveyard files, which contain information about funds that no longer report to a database, to derive approximations of point-in-time available information to counteract survivorship bias. We further address concerns about backfill bias by using hedge fund database inclusion dates to accurately reflect when a specific fund's information became available through either of the two database providers.
[6] This representation of AUM coverage considers the AUM coverage of the hedge fund universe captured by our database in relation to estimates about the overall size of the hedge fund universe from the 3Q 2017 HFR Global Hedge Fund Industry Report.

Exhibit 1: Number of Funds and Total AUM of Hedge Fund Studied Universe

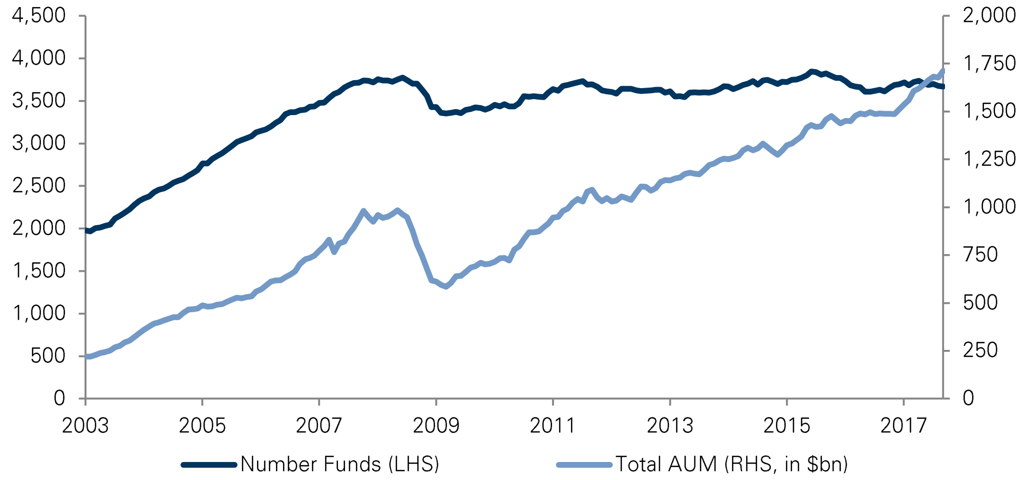

Source: HFR, BarclayHedge, GSAM as of December 2017

In line with the results from Exhibit 1, the average AUM across hedge funds dipped by around $100mn during the Global Financial Crisis in 2008, but has subsequently increased and now sits at two and a half times the level of the average AUM post-Global Financial Crisis (Exhibit 2). While the median AUM generally co-moves with the average AUM, it is worth pointing out that it is noticeably less than $100mn. When contrasting the median and the average, it becomes apparent that the average is skewed by the presence of a few high-AUM funds, which overpowers the presence of a substantial number of smaller AUM funds.

Exhibit 2: Median and Average AUM of Hedge Fund Studied Universe

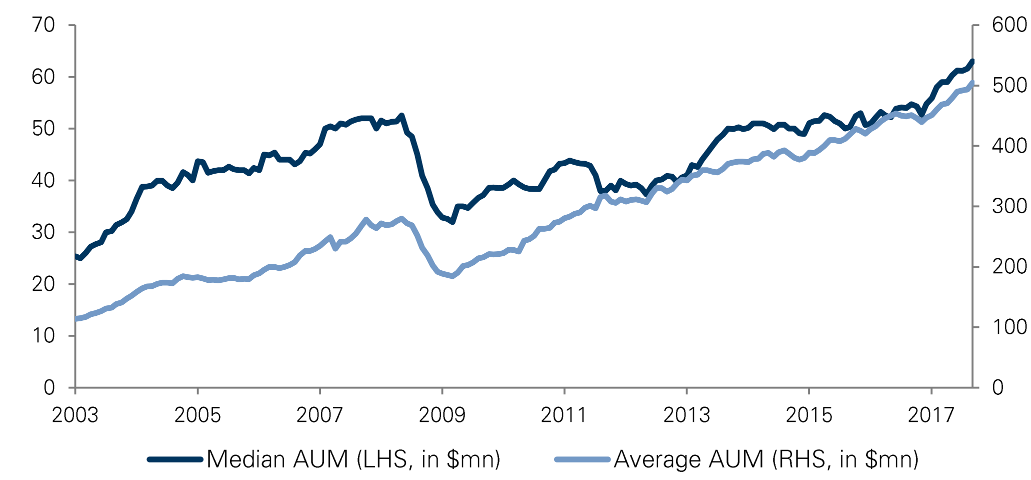

Source: HFR, BarclayHedge, GSAM as of December 2017

When constructing aggregate hedge fund return time series from individual hedge fund information, there are typically two main weighting approaches, AUM-weighting and equal weighting. In contrast to AUM-weighting, equal weighting has the benefit that the composition of the overall hedge fund universe is not dominated by a few very large hedge funds, which is an imminent concern provided the evidence from Exhibit 2.[7] Relatedly, equal weighting implies that our return representations capture all size segments of the hedge fund universe comprehensively. This is particularly relevant in the context of the complexities for hedge fund investors to perform hedge fund due diligence on a large set of hedge funds. Equal weighting has the advantage of providing access to a diverse set of smaller-capitalization funds that investors might otherwise find difficult to subject to a thorough and comprehensive due diligence procedure.

Another key component of the hedge fund universe construction in addition to equal weighting is a "bottom-up" process of grouping hedge funds. Instead of considering the universe of hedge funds as a single abstract average of all available return time series, we break the universe down according to common hedge fund investment styles. These styles represent selections of hedge funds from the overall universe that generally are still broad and diversified, but are more homogeneous than the overall universe in that they share certain investment characteristics. These styles then enable us to develop an understanding of the systematic drivers of their returns, which we subsequently aggregate back to the overall hedge fund universe.

Commonly considered aggregations of hedge fund styles are Equity Long Short, Macro, Relative Value, and Event Driven. Hedge Fund Research, Inc. generally characterizes these four aggregations, which we will refer to as categories, as follows:[8]

(i) Equity Long Short:

This category contains hedge funds, whose exposure - both long and short - is primarily in equities. These funds employ a variety of investment styles, ranging from quantitatively to fundamentally driven approaches.

(ii) Macro:

The Macro category represents funds, whose investment process and resultant exposures to a broad set of different asset classes is predicated on movements in underlying economic variables. Investment theses are based on a variety of discretionary or systematic techniques.

(iii) Relative Value:

Hedge funds in this category take positions across different asset classes in order to exploit valuation discrepancies in the relationship between multiple securities.

(iv) Event Driven:

Hedge funds in this category establish exposures to companies currently or prospectively involved in corporate transactions. The types of such exposures cover the whole spectrum of the corporate capital structure.

Exhibit 3 shows the relative proportions of these categories in December 2017. Equity Long Short hedge funds make up almost 50% of the universe, while Macro hedge funds make up between a quarter and a third. The remainder is split approximately two thirds to one third between Relative Value and Event Driven hedge funds, respectively. This relative composition of the overall universe

[7] Despite concerns about differences in concentration between an AUM-weighted and an equally weighted aggregate hedge fund return time series, it should be noted that these two construction approaches result in fairly highly correlated aggregated return time series. Comparing the HFRI Asset Weighted Index (AUM weighted) to the HFRI Fund Weighted Composite Index (equally weighted) over the maximum available overlapping time period from December 2007 until November 2017, it becomes apparent the two time series are 92.7% correlated with a Root Mean Square Error (RMSE) of 2.3%. For more information about the two hedge fund indices, please refer to the Hedge Fund Research, Inc. website www.hedgefundresearch.com.
[8] Hedge Fund Research, Inc. provides information about hedge fund indices and descriptions of common hedge fund investment styles on their website www.hedgefundresearch.com. The summaries for the four hedge fund categories source information from these descriptions.

does not change much over time. In fact, the average month-on-month change across the weights to all four categories amounts to only slightly below 0.8%.

Exhibit 3: Weighting of Individual Hedge Fund Categories in Studied Universe

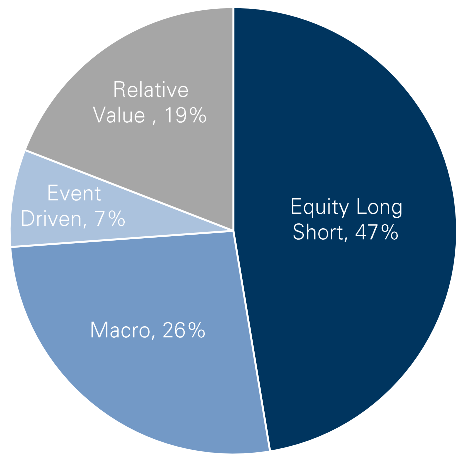

Source: HFR, BarclayHedge, GSAM as of December 2017

2.2 Persistence of Hedge Fund Performance

Having established the hedge fund dataset, we now turn to the analysis of performance persistence. In order to gain a high-level insight into potential performance persistence, we consider return aggregates for the four main categories of the hedge fund universe. Exhibit 4 shows the annual performance of each of these four categories and ranks their performance from 2003 through 2017. While there may have been a certain degree of stability in the very first years of the sample, the ranking of the categories subsequently changes dramatically year over year. The Macro category, for example, jumps from the bottom performer in 2009 and 2010 and again in 2012 and 2013 (0% in each year) to being the second best performer in 2011 (-3%) and even the top performer in 2014 (+6%) before dropping again in 2015 and 2016. Equity Long Short is never the worst performer after 2011 but it alternates year by year between top and third strongest performer. Overall, there is little evidence of performance persistence on this fairly high aggregation level of the four hedge fund categories.

Exhibit 4: Annual Hedge Fund Category Performance

2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017
ED 23%	ED 15%	ELS 12%	ED 14%	ELS 13%	MA 8%	RV 27%	RV 12%	RV 1%	RV 10%	ELS 15%	MA 6%	RV 0%	ED 10%	ELS 14%
ELS 22%	ELS 9%	ED 9%	ELS 14%	MA 11%	RV -20%	ELS 27%	ED 12%	MA -3%	ELS 9%	ED 13%	RV 4%	MA -1%	RV 9%	ED 6%
RV 14%	RV 8%	RV 7%	RV 12%	RV 8%	ED -20%	ED 26%	ELS 11%	ED -4%	ED 8%	RV 8%	ELS 2%	ELS -1%	ELS 5%	RV 5%
MA 11%	MA 2%	MA 5%	MA 7%	ED 7%	ELS -26%	MA 5%	MA 9%	ELS -9%	MA 1%	MA 0%	ED 1%	ED -4%	MA 2%	MA 3%

Legend

ELS – Equity Long Short | MA – Macro | RV – Relative Value | ED – Event Driven

Source: HFR, BarclayHedge, GSAM as of December 2017

In order to more accurately reflect the challenges in assembling a hedge fund portfolio, we complement this high-level analysis with a fund-level analysis of persistence. Corresponding results in the academic literature are mixed. Agarwal and Naik (2000a), Agarwal and Naik (2000b), Amenc et al. (2003) as well as Bares et al. (2003), for example, have established evidence in favor of performance persistence for shorter periods up to a quarter. Ter Horst and Verbeek (2007), Boyson (2008) and Eling (2009) provide a more nuanced perspective that is supportive of performance persistence for shorter-term periods up to six months, but regard the evidence for intermediate- to longer-term horizons as more challenged. These intermediate- to longer-term results are in line with Brown and Goetzmann (2003), Capocci and Huebner (2004), Capocci et al (2005) and Malkiel and Saha (2005).

Acknowledging the practical complexities in adjusting hedge fund portfolios dynamically, this analysis focuses on an annual period to evaluate performance persistence in single hedge funds. For each year from 2003 until 2015, we sort all hedge funds that have reported returns throughout the entire year into performance quintiles. We subsequently measure the performance over the following year and apply another quintile sort. For the following year's performance, we however need to be mindful that hedge funds may no longer report returns to the hedge fund database providers. This may be driven by, for example, fund restructurings or liquidations. For this reason, the ranking in the subsequent year also contains a column termed "NR", which stands for "Not Reporting." This column reflects those funds that have stopped reporting returns at some point throughout the subsequent year.

Exhibit 5: Transition Matrix for Performance Quintiles of Individual Hedge Funds

		Subsequent Year Ranking					
		1	2	3	4	5	NR
Initial Year Ranking	1	29%	18%	13%	12%	20%	8%
	2	17%	21%	19%	15%	15%	12%
	3	13%	18%	19%	18%	14%	18%
	4	12%	16%	15%	17%	15%	25%
	5	18%	11%	10%	12%	22%	27%

Source: HFR, BarclayHedge, GSAM as of December 2017

Exhibit 5 contains 55,727 observations from 2003 until 2016. For each row, the different columns show how likely a fund is to end up in the respective performance quintiles in the following year.[9] For example, for a fund that is initially ranked in the third quintile, there is a 13% likelihood that it will be in the first quintile in the subsequent year and an 18% likelihood that it will be in the second quintile. High performance persistence would be demonstrated by the diagonal elements of this matrix being an order of magnitude larger than the off-diagonal elements. While we find some very limited evidence for this effect for the very best and worst performing hedge funds in the initial year ranking, instability abounds and one even observes evidence of mean reversal of returns in the extreme quintiles.

As a matter of fact, the probability of starting off in quintile 1 and ending in the worst performing quintile is the second highest probability after staying in quintile 1. The same holds true for the worst performer where moving from quintile 5 to quintile 1 in the following year has the second highest probability after remaining at the bottom. Generally, in contrast to the required pattern to establish performance persistence, each row in Exhibit 5 actually displays a much more pronounced tendency towards a uniform distribution of likelihoods across the different quintiles. Overall, Exhibit 5 confirms the lack of unified evidence in the academic literature of performance persistence in single hedge funds once one imposes a minimum evaluation time period.

Another point to note about Exhibit 5, which is problematic for the selection of portfolios consisting of only a few individual funds, is the high likelihood of a fund not reporting 12 months of returns in the subsequent year. While there is already an approximately 1 in 13 likelihood that funds in the top quintile do not report returns in the following year, this probability increases monotonically for worse-performing quintiles and exceeds a 1 in 4 likelihood for the worst performing quintile. It is noteworthy that these likelihoods only represent one-year quantities and imply an even higher fraction of hedge funds that may potentially stop reporting over a multi-year period.[10]

This type of inevitable hedge fund turnover may lead to potentially costly searches for replacement funds and may involve periods where certain fractions of a hedge fund portfolio are left unallocated and therefore cannot deliver the return characteristics that investors seek. This is a challenge to which the proposed alternative approach to hedge fund investing will not be subject.

2.3 Convergence of Hedge Fund Portfolios to the Overall Hedge Fund Universe

While the lack of performance persistence warrants caution when it comes to the construction of select hedge fund portfolios, the question arises whether selections of hedge funds could provide sufficient diversification to deliver alternative returns without the risk of exposing a portfolio to the idiosyncrasies of individual hedge funds while still offering the potential to generate superior risk-adjusted returns.

Exhibit 6 provides answers to this question by comparing portfolios of differing number of hedge funds to a broad universe of hedge funds as well as to the average performance of funds in that universe. In this analysis we randomly form hedge fund portfolios of various sizes and hold these portfolios for a period of five years using data covering a time period from October 2012 to September 2017.[11] The portfolio sizes we consider range between 5 and 200 funds. We then run a

[9] Each element in this matrix is the average over the transition likelihoods for all initial sorts from 2003 and 2015.
[10] For example, if a hedge fund starts off in quintile 1, there is approximately a 36% likelihood that this fund will stop reporting at some point in the subsequent 3 years.
[11] If a hedge fund ceases to publish returns during the time frame considered for this analysis, we re-allocate its weight to the remaining hedge funds in the respective sampled portfolios of hedge funds. If all hedge funds from

bootstrapping analysis of 10,000 selections per portfolio size and calculate the Sharpe ratio as well as the correlation to the average return of all hedge funds in our database for the analyzed time period, for each random selection.

Exhibit 6 displays the Sharpe ratio and correlation characteristics for the distribution associated with each specific hedge fund portfolio size in the simulation. The most striking feature of this analysis is the high correlation of the simulated portfolios with the average returns across all hedge funds in our universe. For a portfolio with only 5 member hedge funds the correlation is at 0.45 and increases to 0.69 for a portfolio of 20 hedge funds. This illustrates how even portfolios with a relatively small number of hedge funds behave very similar to the average return across all hedge funds. The average Sharpe ratios of the simulated portfolios are below the ones from the hedge fund average returns but converge for larger portfolios. This is partially driven by the diversification effect of larger portfolios given that the applied selection mechanism does not model any skill in selecting hedge funds. However, the dispersion between the 5th and 95th percentile illustrates the variability in terms of Sharpe ratio that the simulation is still subject to across different portfolio sizes.

Exhibit 6: Sharpe Ratio and Correlation for Simulated Hedge Fund Portfolios (October 2012 to September 2017[12])

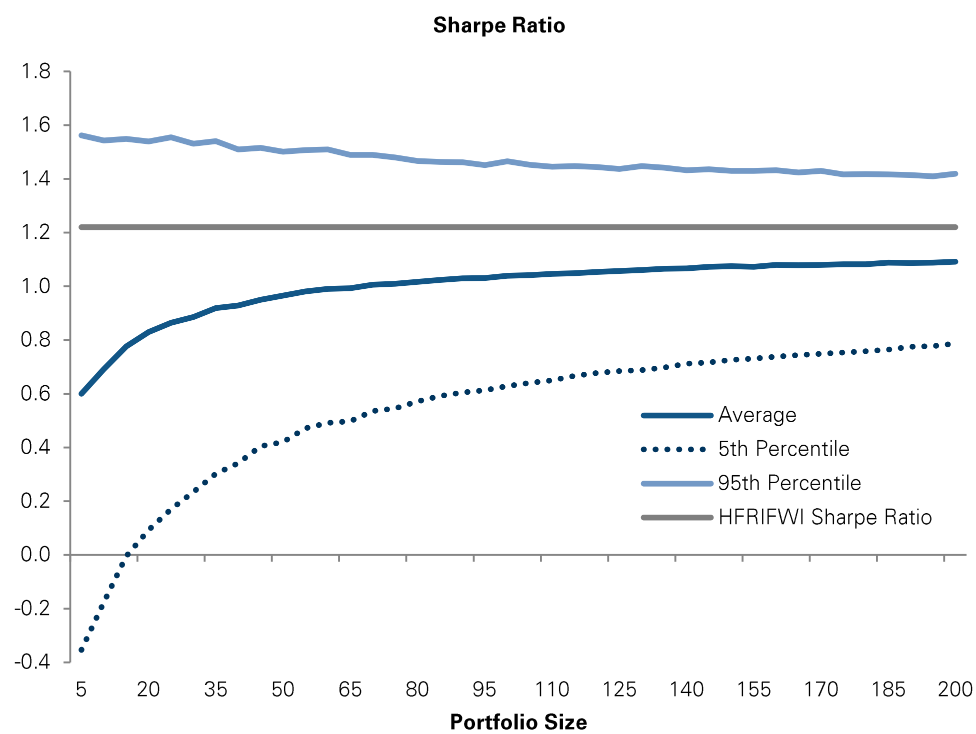

an initial selection cease to publish returns, we there are no more hedge funds in an initial selection from the universe

[12] A five year time period is used for the simulation in order to ensure the inclusion of an appropriate number of Funds with overlapping time periods without inducing excessive survivorship bias in the analysis.

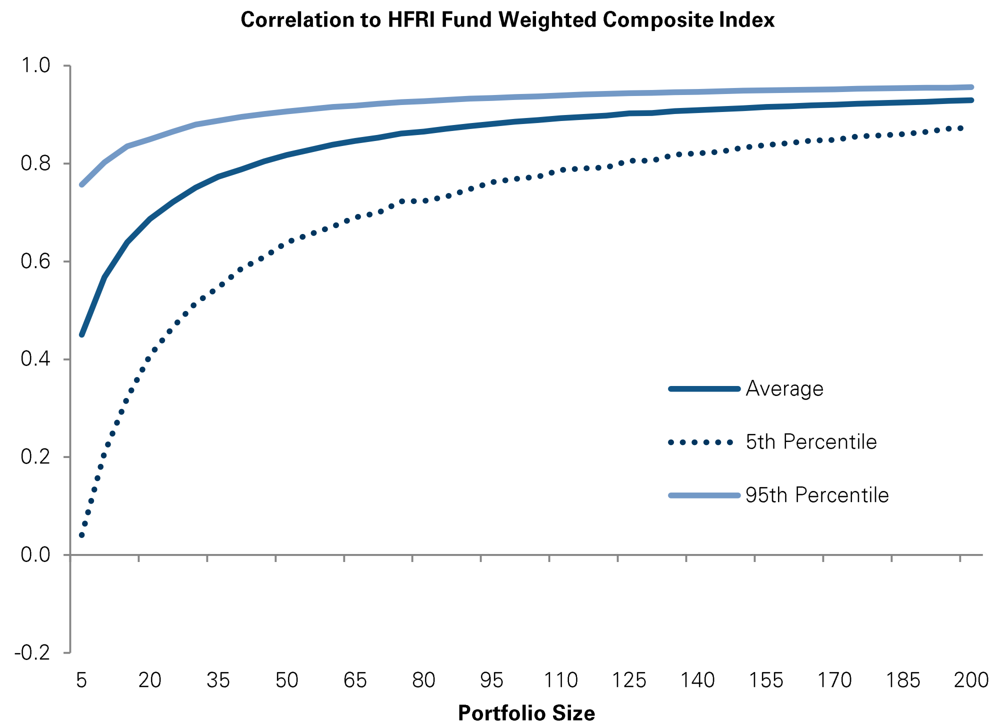

Source: HFR, BarclayHedge, GSAM as of December 2017

While the simulation study relies on indiscriminate selections of hedge funds, the following paragraphs complement this analysis by analyzing portfolios of fund-of-funds, which deliberately select specific funds from the hedge fund universe that they cover. Exhibit 7 highlights the risk-return characteristics of our overall representation of the universe of individual hedge funds in conjunction with the characteristics for a universe of fund-of-funds. It considers single hedge fund as well as fund-of-fund data[13] over the past five years up until the fourth quarter of 2017 and also includes the performance of the equally weighted average return across all hedge funds over the same time period.

It becomes apparent that fund-of-funds generally accomplish diversification, as their distribution is located within the distribution of the overall universe of hedge funds. The average volatility across all fund-of-funds is approximately 6.1%, while the average for all individual hedge funds amounts to close to 11.4%. It is however not necessarily the case that the additional diversification translates into superior risk-adjusted returns, particularly when compared to a diversified aggregate of individual hedge funds. The average time series constructed from the universe of all hedge funds delivers a return of 5.2%; more than 1% higher than the 4.1% return of the fund-of-fund universe at a risk level of around 3.7%, which is 2.4% lower than the average volatility of the fund of funds at 6.1%. This translates into a Sharpe ratio of 1.3 for the average time series constructed from the universe of all hedge funds, which is 0.4 higher than the average Sharpe ratio of the fund-of-funds (0.9).

[13] We construct a universe of fund of funds analogous to the construction of the universe of single hedge funds as outlined in Section 2.1.

Exhibit 7: Historical Risk/Return Distribution

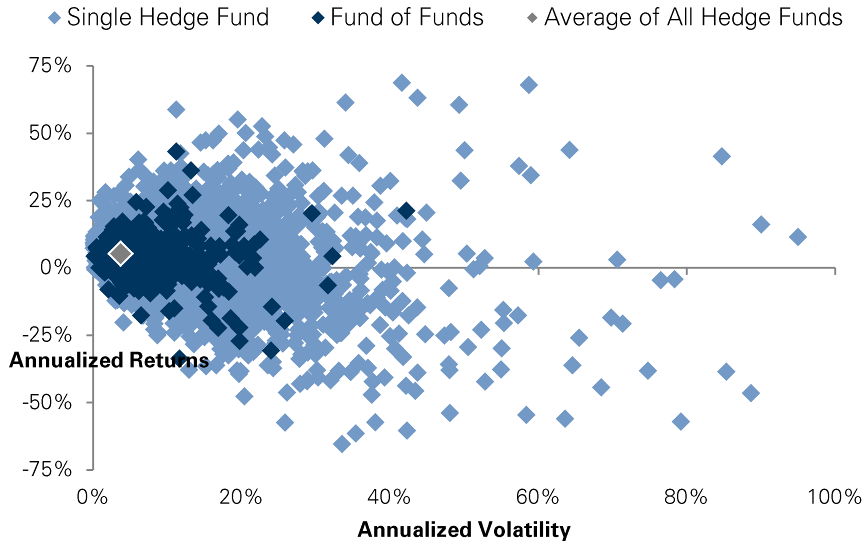

Source: HFR, BarclayHedge, GSAM as of December 2017

In summary, the lack of performance persistence and its implications for the necessary turnover of investors' hedge fund portfolios may provide an argument against hedge fund portfolios with very few individual funds. If a hedge fund investor deviates from a very select portfolio by increasing the number of funds, the resulting performance may already exhibit a high degree of resemblance, on average, with a broad and diversified set of hedge funds. However there is still substantial risk to deviate from the broader universe, as evident from the deviation in Sharpe ratios in the top chart of Exhibit 6.[14] Seeking exposures of a broadly diversified portfolio of hedge funds instead is an effective means for a hedge fund investor to navigate this risk. Such portfolio furthermore proves to exhibit attractive risk-adjusted return characteristics, even compared to the average fund-of-funds, as highlighted in Exhibit 7.

3 Systematic Drivers of Hedge Fund Performance

Building on a broadly diversified portfolio of hedge funds, our approach to identifying the systematic risk exposures delivered by hedge funds consists of three steps: First, the identification of our universe of hedge fund returns together with a hedge fund categorization scheme. As discussed in Section 2.1, we break the universe down into four main categories. Within each category, we then identify individual hedge fund styles, for which we aim to characterize the systematic return drivers. The second step is the identification of a selection of factors which can be classified as either traditional or alternative risk premia associated with each of the different hedge fund styles within a hedge fund category. Finally, these two steps are tied together by a weight estimation methodology, which is applied for each hedge fund style and determines exposures to traditional and alternative risk premia in order to best emulate a given hedge fund style's returns.

[14] In unreported results we repeat the simulation analysis using the information ratio versus the MSCI World index rather than the Sharpe ratio as performance metric. The results are similar in as much as the random portfolios converge to the information ratio of the average returns across all hedge funds as the portfolio size grows. The one noteworthy difference is that the information ratio decays as the portfolio size increases (for example, from an average of 0 for a portfolio of 10 holdings to an average of -0.03 for 200 holdings).

3.1 Characteristics of Systematic Factors in Hedge Funds

The approach used to identify systematic factors delivered by hedge funds is based on insights from the academic literature on common risk premia for mutual funds. The advent of factor analysis of mutual fund returns can be traced back to the Capital Asset Pricing Model (CAPM) by Sharpe (1964), Lintner (1965), and Mossin (1966) who link return expectations back to exposures to the equity market factor. Fama and French (1992, 1993) extend this factor set by a value and a size factor and apply the resulting 3-factor model to equity returns. Carhart (1997), based on Jegadeesh and Titman (1993), extends the Fama and French (1992, 1993) factors by a momentum factor and finds that there is a significant loading on this factor in the cross-section of mutual funds.

Based on this work on factor analysis for mutual funds, Fung and Hsieh (1997) pioneer the analysis of the systematic return drivers for hedge fund styles. Research by Schneeweis and Spurgin (1998), Liang (1999), Edwards and Caglayan (2001), Capocci and Hubner (2004) as well as Hill et al (2004) refines the factor set used to determine the drivers of the returns of hedge fund styles by focusing on more easily interpretable factors as well as by considerations around tradability. Fung and Hsieh (2001, 2004) as well as Agarwal and Naik (2000a, 2000b, 2004) further expand the set of return drivers beyond the inclusion of basic representations of asset classes or parts thereof by introducing implementable trading strategies to improve the explanatory power of their approximation of hedge fund returns. Their factor sets can already be decomposed into traditional and alternative risk premia, with both types of factors playing a key role in understanding and emulating the risk-return characteristics of hedge funds.

As defined in the introduction, traditional risk premia are individual "long only" market factors (betas) such as equities or fixed income. Alternative risk premia instead are systematic, multi-asset, long/short investment strategies, backed by academic research and employed by market practitioners. Roughly, alternative risk premia fall into four categories:

(i) Value strategies, which take advantage of the tendency for cheap assets to outperform expensive assets on a relative basis;

(ii) Carry strategies, which capitalize on the tendency for higher yielding assets to outperform lower yielding assets;

(iii) Momentum strategies, which exploit the tendency for recent relative price movements to continue in the near future; and

(iv) Structural strategies, which capture returns from market anomalies arising from structural constraints rather than economic fundamentals.

Attractive risk-adjusted returns, return persistence, economic intuition, and their highly liquid and cost-efficient profile have led an increasing number of investors to adopt alternative risk premia strategies in their portfolios. Many such strategies have historically realized low correlation to the price movements of traditional asset classes, and have proven effective in explaining sizeable portions of the returns of particular hedge fund styles.

3.2 Mapping Systematic Factors to Hedge Fund Categories

As outlined in Section 2.1, we do not just consider a single representation of the hedge fund universe as a whole, but we rather aim to develop a precise and tailored understanding of the traditional and alternative risk premia factors that play a role for each hedge fund category. The applied approach to factor identification even goes a level deeper to not only look at individual hedge fund categories but to consider aggregates of hedge funds within a category, so-called styles, that share commonalities in terms of the investment approach as well as investment exposures.

When identifying the appropriate factor set for specific styles within an individual hedge fund category, we rely on fundamental analysis verified by a quantitatively driven weight estimation methodology. Fundamental insights allow us to cross-validate factors using a range of qualitative sources from hedge fund database information to prime brokerage reports, hedge fund consultant reports or hedge fund holdings from 13F filings.[15] This approach puts us in a position to not only identify correlation between hedge funds and risk premia, but also to address causation, which is beneficial for the out-of-sample properties that the estimated weights will exhibit to the returns of the hedge fund style under consideration.

The following overview outlines general characteristics for the identification and selection of traditional and alternative risk premia. For ease of presentation, the overview aggregates these characteristics to the level of the four main hedge fund categories identified in Section 2.1:

(i) Equity Long Short:

A core exposure of funds within the Equity Long Short category is global equity market exposure. This can be complemented by additional traditional risk premia providing exposure to equity sectors actively held by Equity Long Short funds, such as energy, technology, or health care. Alternative risk premia such as Value strategies further complete the set of exposures. Finally, systematic stock selection aspects can be captured with a factor based on 13F filings.

(ii) Macro:

The core exposures for this hedge fund category are alternative risk premia – specifically Momentum strategies across a diverse set of asset classes. From the perspective of alternative risk premia, Carry strategies in foreign exchange also contribute to understanding the drivers of Macro hedge fund returns. Traditional risk premia representing exposures to, for example, commodities or emerging market equities exhibit a suitable degree of complementarity to the aforementioned alternative risk premia.

(iii) Relative Value:

Risk exposures for the Relative Value category consist of a diverse set of traditional risk premia paired with alternative risk premia falling into the category of Structural strategies. The set of traditional risk premia is fairly diverse in this hedge fund category, consisting of not only exposures at various seniority points of the corporate balance sheet, but also of government debt instruments, Master Limited Partnerships (MLPs) as well as Real Estate Investment Trusts (REITs). With respect to alternative risk premia, factors with return profiles similar to those of illiquid strategies[16] arise from index option strategies as well as from the optionality component in convertible bonds.

[15] 13F filings refer to Form 13F by the US Securities and Exchange Commission (SEC). Institutional investment managers satisfying certain criteria such as holding more than $100mn in qualifying assets need to submit this form on a quarterly basis. The form contains information about the holdings of those investment managers. Filings are made publicly available with a 45-day delay after the end of each calendar quarter. See the SEC website https://www.sec.gov/divisions/investment/13ffaq.htm for more information.
[16] These can be broadly understood as patterns of smooth accumulation of performance with intermittent periods of sharp drawdowns.

(iv) Event Driven:

Similar to the Relative Value category, alternative risk premia exposures capture illiquidity-type return profiles and fall into the Structural strategies block. The set of traditional risk premia provides exposure to different levels of market capitalization for equities as well as to different seniority points of the corporate balance sheet.

3.3 Principles of the Weight Estimation Methodology

The weight estimation builds on original insights from Sharpe (1992), who uses factors to decompose and understand the returns of mutual funds and suggests a framework which actually invests in the respective factors in order to mimic mutual fund returns. More explicitly, we lean on subsequently developed approaches proposed, for example, by Hasanhodzic and Lo (2007), Hill et al (2004), and Jaeger (2008), which take Sharpe's (1992) methodology and extend it further to mimic the returns of hedge funds.

For a hedge fund style S, the methodology focuses on error terms of the form

$$e^{S}(\mathrm{t}) = HFS^{S}(\mathrm{t}) - \left[\alpha + \sum_{\mathrm{f}\in TRP^{S}} \beta_{\mathrm{f}} \cdot \mathrm{RF}_{\mathrm{f}}(\mathrm{t}) + \sum_{\mathrm{f}\in ARP^{S}} \beta_{\mathrm{f}} \cdot \mathrm{RF}_{\mathrm{f}}(\mathrm{t})\right],$$

where α denotes a constant and β_{f} denotes the weight estimate for a risk premium from either the set of "Traditional Risk Premia" $\mathrm{TRP}^{\mathrm{S}}$ or the set of "Alternative Risk Premia" $\mathrm{ARP}^{\mathrm{S}}$ for hedge fund style S. We further denote the excess return of the average return of hedge fund style S by $\mathrm{HFS}^{\mathrm{S}}$ and the excess return of risk premium f (traditional or alternative) by RF_{f}. As outlined in Section 2.1, $\mathrm{HFS}^{\mathrm{S}}$ represents an equally weighted average of the returns of a subset of hedge funds from the hedge fund dataset that we construct from single hedge fund time series originally provided by the two data providers. For t spanning a time period of 24 months, we then determine α and the vectors $\{\beta_{\mathrm{f}}\}_{\mathrm{f}\in\mathrm{TRP}^{\mathrm{S}}}$ and $\{\beta_{\mathrm{f}}\}_{\mathrm{f}\in\mathrm{ARP}^{\mathrm{S}}}$ that minimize a quadratic transformation of the error terms $\mathrm{e}^{\mathrm{S}}(\mathrm{t})$.[17]

In line with the philosophy of the academic literature that originated from Sharpe (1992), we aim to translate the outcome of the in-sample weight estimation methodology into an out-of-sample portfolio allocation, which is the core of our construction of liquid representations of the factor exposures that hedge funds exhibit. As outlined in the introduction, we generally distinguish between traditional and alternative risk premia exposures for hedge funds and acknowledge the existence of an unexplained portion. The first two are incorporated in the factor sets $\mathrm{TRP}^{\mathrm{S}}$ and $\mathrm{ARP}^{\mathrm{S}}$, specified for each hedge fund style S, and what we refer to as the unexplained portion is captured by the constant term α. This portion, for example, reflects the fact that, by construction, the well-defined and liquid risk premia may naturally exhibit a degree of divergence to the opaqueness and illiquidity of some hedge fund investment strategies. It is then crucial for the determination of the overall success of the weight estimation procedures to verify the relative proportions of unexplained returns and returns driven by the two classes of risk premia, which we will further elaborate on in the subsequent section.

For the out-of-sample implementation of the methodology, this implies in turn that only the components from the factor sets $\mathrm{TRP}^{\mathrm{S}}$ and $\mathrm{ARP}^{\mathrm{S}}$ can be made available to an investor, as

[17] The discussed transformation creates a convex objective function that ensures that the minimization problem is well-defined. It overweights more recent observation and also controls for illiquidity-induced autocorrelation using an adaptation of the methodology proposed by Scholes and Williams (1997). Note that the objective function is dynamic in the sense that it will change each month based on updated data points, albeit that the actual transformation function is static.

$$\sum_{\mathrm{f} \in TRP^S} \beta_{\mathrm{f}} \cdot \mathrm{RF}_{\mathrm{f}}(\mathrm{u}) + \sum_{\mathrm{f} \in ARP^S} \beta_{\mathrm{f}} \cdot \mathrm{RF}_{\mathrm{f}}(\mathrm{u}),$$

where u denotes an out-of-sample time period that occurs after any of the periods t used for the in-sample weight estimation. As this process relies on collated hedge fund data, there is an inevitable gap between u and any of the respective periods t in order to account for the publication lag inherent in any hedge fund database. Once this delay has passed, weights to sets of traditional and alternative risk premia factors are re-estimated on a monthly basis based on the most recently available hedge fund database information by both data providers. This monthly re-estimation of weights is targeted towards capturing the dynamic nature of hedge fund positioning. It complements the other source of dynamism present in this portfolio construction, which arises from shifts in investment exposures within each of the alternative risk premia.

The next crucial step is the aggregation to the level of the overall hedge fund universe. Even though we identify sets of traditional and alternative risk premia for individual styles within hedge fund categories, the objective remains to provide access to the return profile of the overall hedge fund universe. We accomplish this aggregation by weighting sets of estimates for traditional and risk premia by the relative number of funds captured within a specific style, in line with the equal weighting approach outlined in Section 2.1.

The out-of-sample implementation can be further adjusted to make the performance more realistic from the point of view of an investor. First, this entails certain assumptions about the trading costs that the implementation of the portfolio of traditional and alternative risk premia might incur in the marketplace. Second, we will also assume a hypothetical management fee of 75bps that an investor might face. The final net performance of the portfolio of traditional and alternative risk premia is what we will refer to as "Liquid Tracking Portfolio" below. It can then be compared to the performance of the average of returns across the broad and diversified universe of hedge funds, as described in Section 2.1, referred to as "Hedge Fund Index." It is important to note that, while the Liquid Tracking portfolio is explicitly tradable, the Hedge Fund Index is merely a representation of average hedge fund performance that is not actually investable and therefore directly accessible to investors. This non-tradability mainly arises because of the sheer scope of the universe covered as well as liquidity and turnover restrictions that investors face in emulating the composition of the aggregate hedge fund universe.

A final noteworthy aspect of the applied weight estimation methodology is its linearity. This paradigm is, for example, challenged by Kat and Palaro (2005) as well as Amenc et al (2008, 2010), who suggest non-linear regression approaches as well as distribution-based considerations. Hasanhodzic and Lo (2007) and Bollen and Fisher (2014) however counter their suggested enhancements in favor of a linear relationship. Besides the case for simplicity in the identification mechanism as well as in the translation of the in-sample estimation to the out-of-sample portfolio of factors, their argument rests on the preferable out-of-sample performance of linear approaches compared to non-linear approaches that tend to be prone to overfitting. Furthermore, distribution-based approaches only match the distribution characteristics in the longer term, which could lead to substantial return mismatches over shorter periods of time.

4 Liquid Tracking Portfolio Simulated Performance

4.1 Performance Comparison

This section reviews the simulated performance of the Liquid Tracking Portfolio whose construction we described in the previous section. We initially discuss the simulated performance of the Liquid Tracking portfolio relative to the Hedge Fund Index before switching the focus to an attribution analysis of the overall hedge fund universe. Both in terms of return contribution as well as marginal contribution to risk, this enables us to explicitly assess the fraction of hedge fund performance that is due to traditional and alternative risk premia and compare it to the fraction that is left unexplained.

Exhibit 8 compares the performance of the Hedge Fund Index to the simulated performance of the Liquid Tracking Portfolio for a period of almost 15 years.[18] The Liquid Tracking Portfolio delivers an annualized simulated return that only falls 1% short of that of the hedge fund index, which translates into a Sharpe ratio difference of less than 0.1.[19] As the Liquid Tracking Portfolio is constrained by construction, as outlined in Section 3.3, to exclude the unexplained part of the returns of the hedge fund universe, we expect the volatility of the Liquid Tracking Portfolio to be below that of the hedge fund index, as the unexplained return component will, by definition, be uncorrelated with the liquid and alternative risk premia but has itself non-negligible volatility. This is confirmed by Exhibit 8.

In terms of co-movement between the two time series, the Liquid Tracking Portfolio exhibits a monthly return correlation of 93.5% to the Hedge Fund Index, i.e. the return observations of the liquid tracking align well with those of the Hedge Fund Index. The close co-movement not only in direction but also in quantity is further substantiated by an annualized Root Mean Square Error (RMSE)[20] of 2.1%.

Exhibit 8: Aggregate Performance Comparison of Hedge Fund Index and Liquid Tracking Portfolio

April 2003 – September 2017	Hedge Fund Index	Liquid Tracking Portfolio (Simulated)
Total Return (Annualized)	6.2%	5.2%
Volatility (Annualized)	5.8%	4.9%
Sharpe	0.83	0.77
Maximum Drawdown	-18.1%	-14.0%
Correlation	-	93.5%
RMSE (Annualized)	-	2.1%

Source: HFR, BarclayHedge, GSAM as of December 2017. As outlined in Section 3.3, the Liquid Tracking Portfolio is net of assumed transaction costs and 75bps management fee.

The high degree of co-movement is driven by a high in-sample quality of fit of the weight estimation procedure, which carries over to the out-of-sample performance displayed in Exhibits 8 and 9. This provides evidence for the appropriateness of the concept of relying on historic weight estimates to determine forward-looking risk exposures that we posit for the hedge fund index in the out-of-sample performance analysis. An approach like this necessitates that the turnover of the weight estimates is

[18] The time window for this analysis is curtailed by the availability of the time series for some of the alternative risk premia.
[19] Note that, according to the single hedge fund assessment of unexplained returns from Section 3.3, only less than half of the hedge funds actually drive the outperformance of the overall hedge fund universe.
[20] The Root Mean Square Error represents the square root of the average squared difference between predicted values (here: Liquid Tracking Portfolio) and observed values (here: Hedge Fund Index).

limited, which is confirmed by an average monthly turnover of the weight estimates of 7.3% (with a standard deviation of 3.7%) for the Liquid Tracking Portfolio.

At the same time, the turnover figures provide evidence for a certain degree of adaptability in the weight estimation methodology. Necessarily, the process needs to be able to detect and react to shifts in the role of certain risk premia (traditional or alternative) over time. For example, Cai and Liang (2012) and Patton and Ramadorai (2013) confirm this notion and emphasize the varying nature of exposures hedge funds take and the need to have the ability to react to such changes. Our methodology accomplishes this objective through turnover in the weights estimated for individual risk premia as well as through allocation changes in the investment strategies inside individual alternative risk premia strategies. For example, in the portion of the Liquid Tracking Portfolio capturing Macro hedge funds, the month-on-month turnover of the risk premia weights is around 5%, while the turnover within the alternative risk premia used for Macro hedge fund tracking can be much higher, as illustrated by a month-on-month turnover of 290%[21] for the Momentum strategies employed in this category.

Exhibit 9: Time Series Performance Comparison of Hedge Fund Index and Liquid Tracking Portfolio

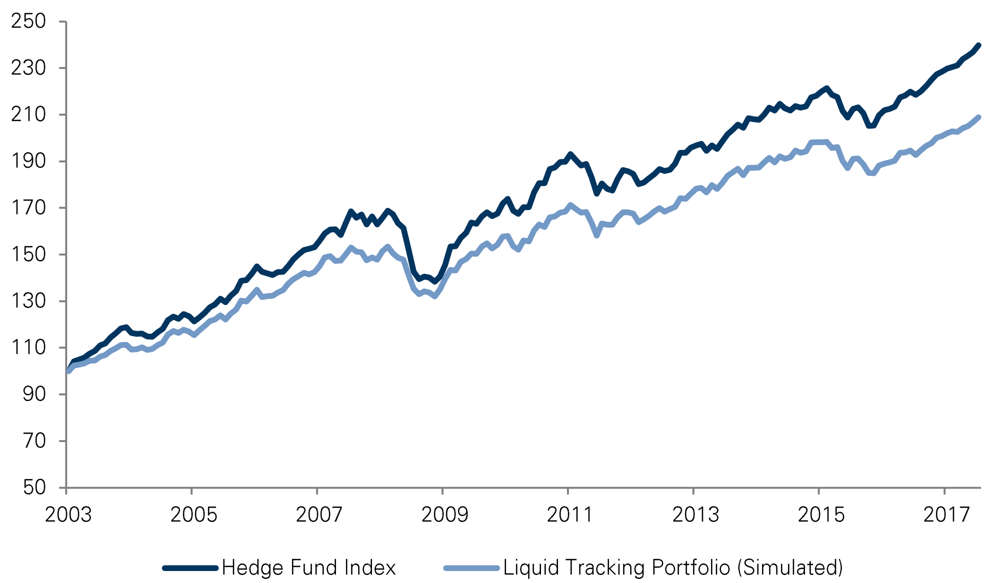

Source: HFR, BarclayHedge, GSAM as of December 2017. As outlined in Section 3.3, the Liquid Tracking Portfolio is net of assumed transaction costs and 75bps management fee.

While Exhibit 8 presents aggregate statistics for the Hedge Fund Index and the Liquid Tracking Portfolio, Exhibit 9 shows the evolution of both time series. It is apparent that the degree of co-movement between the two time series is very consistent over time and that there are no periods of significant divergence.

While the co-movement is very consistent in the time series representation of Exhibit 9, it turns out that there is a sizeable degree of cross-sectional variation in how the 1% annualized performance difference of the Hedge Fund Index to the Liquid Tracking Portfolio is distributed among hedge funds. Based on the style-by-style portfolio construction of selected traditional and alternative risk premia outlined in Sections 3.2 and 3.3, one can construct performance comparables for individual

[21] In order to put the turnover figures in context, over the time frame in question, the Macro Liquid Tracking portfolio and the Momentum Alternative Risk Premium realized annualized volatilities of 4.5% and 9.4%, respectively.

hedge funds according to the hedge fund style that each hedge fund is categorized in.[22] This way, although the general focus lies on the aggregate hedge fund universe, it is possible to make inferences about the cross-sectional distribution of the unexplained returns in the overall universe of hedge funds.

For the hedge fund sample outlined in Section 2.1, which covers a period of almost 15 years with initially around 2,000 funds that later grows to close to 4,000 funds, it turns out that only 45.8% of the funds actually manage to have positive unexplained returns when measured against their liquid performance comparable. At the same time however, there is a considerable degree of variation in the unexplained returns. According to our analysis, while the 75th percentile of hedge funds manages to realize 47bps of monthly positive unexplained performance, the 25th percentile falls short by 75bps per month. Keeping in mind the 1% overall performance difference between the Hedge Fund Index and the simulated Liquid Tracking Portfolio, this points to a fairly high degree of concentration of positive unexplained returns within the universe of hedge funds. This consideration reiterates difficulties hedge fund investors may face in their allocation to individual funds.

4.2 Decomposition of Hedge Fund Performance

On the aggregate hedge fund universe level, the previous section demonstrates the co-movement between the Hedge Fund Index and the out-of-sample performance represented by the simulated Liquid Tracking Portfolio. Below, we will quantify the return and risk contributions of the unexplained returns of the Hedge Fund Index relative to the proposed Liquid Tracking Portfolio and put them into comparison with the impact of the traditional and alternative risk premia.[23]

Exhibit 10 presents the return contribution[24] as well as the marginal contribution to risk of the returns of the Hedge Fund Index coming from unexplained returns and traditional and alternative risk premia. In line with Exhibits 8 and 9, the fraction of returns attributed to unexplained returns is only 16% of the overall returns of the Hedge Fund Index, with the remaining portion of 84% attributable to traditional and alternative risk premia. Further breaking down the return contribution of the two classes of risk premia, the return split between traditional and alternative risk premia comes out at approximately 55/45, which is a clear indication of the important and sizeable contribution that alternative risk premia make towards capturing hedge fund returns.

[22] We compare the cumulative performance of each hedge fund captured by the analysis over all months that this fund has a return observation in our database to the performance of the liquid portfolio of traditional and alternative risk premia constructed for the hedge fund style, under which the specific hedge fund falls, over the same months.
[23] In an out-of-sample context, unexplained returns can essentially be decomposed into two parts: (1) Unexplained returns from the in-sample weight estimation procedure and (2) prediction error arising from the process of inferring out-of-sample weights from in-sample estimates. The prediction error in (2) can further be decomposed into a portion that arises as exposures to traditional and alternative risk premia change during the out-of-sample period compared to the window used for estimation as well as a portion attributable to the relative proportions of the different hedge fund investment styles changing over time. While the effect of changing weights has already been addressed in the context of the discussion about turnover in the previous section, it also turns out that the relative weight shifts of individual styles are minor, in line with the evidence presented in Exhibit 3 in Section 2.1 for the four main hedge fund categories.
[24] While Exhibit 8 presents annualized total returns, the return decomposition in Exhibit 10 uses non-annualized return quantities. 136.9% total return over the time period considered translates to 6.2% annualized total return.

Exhibit 10: Factor Attribution of Hedge Fund Index Performance

April 2003 – September 2017	Relative Return Contribution	Relative Marginal Volatility Contribution
Unexplained	15.9%	20.3%
Alternative Risk Premia	37.5%	13.6%
Traditional Risk Premia	46.5%	66.1%
Aggregate	100.0%	100.0%

Source: HFR, BarclayHedge, GSAM as of December 2017

In terms of marginal contribution to risk, the breakdown between traditional and alternative risk premia shifts towards traditional risk premia which explain about 66% of the overall volatility. This is driven by the directional nature of the traditional risk premia, which tends to imply higher volatility for these factors, in comparison to the more diversified and long/short types of exposures typically embodied by alternative risk premia. For the unexplained return component, the contribution to the overall volatility remains at a level (~20%) that is similar in magnitude to the proportional contribution to returns.

An important determinant of the stability of the out-of-sample contribution analysis in Exhibit 10 is the complementarity of the individual components of the return and risk contribution breakdown. For this reason, Exhibit 11 displays the pairwise correlations of the three hedge fund return components. Since the unexplained portion of the returns is orthogonal to traditional and alternative risk premia, we expect the correlation of unexplained returns to the other factors to be close to zero, which is confirmed for alternative risk premia and to a lesser degree for the traditional risk premia. We attribute the residual correlation between unexplained returns and traditional risk premia to short-term market timing by some hedge fund styles, which only get picked up in an uncomplete manner by the monthly weight estimation process.

Exhibit 11: Correlation of Hedge Fund Attribution Factors

April 2003 – September 2017	Unexplained	Alternative Risk Premia	Traditional Risk Premia
Unexplained	100%	-0.3%	25.4%
Alternative Risk Premia		100%	17.7%
Traditional Risk Premia			100%

Source: HFR, BarclayHedge, GSAM as of December 2017

A final point to highlight about Exhibit 11 is the low correlation between traditional and alternative risk premia. This bodes well not only for the stability of the contribution analysis, but also highlights the complementarity of the role that alternative risk premia play in explaining hedge fund returns out-of-sample in the applied methodology over and above the attribution that can already be inferred from traditional risk premia.[25]

Exhibit 12 elaborates further on the return decomposition from Exhibit 10 by breaking down the contribution into three subperiods of approximately five years each. At first, it is noteworthy that hedge fund performance has actually undergone quite a high degree of time variation, as evidenced by the aggregate of the three columns displayed for each time period. A period of exceptionally strong returns in the run-up to the Global Financial Crisis is followed by a period of more challenged performance thereafter, which has then given way to a slight performance improvement in the latest part of the sample. Assessing the impact of the individual components, the exhibit proves the

[25] The low residual correlation is predominantly driven by the Momentum strategies present in the Macro category that can take directional exposures based on sustained price moves in assets that also reflect traditional risk premia.

consistency of the return contribution of the alternative risk premia, as the impact of alternative risk premia has a higher contribution than that of unexplained returns in each of three subperiods.

In terms of the relative contribution of traditional and alternative risk premia, it becomes apparent that the 55/45 overall split is similar in the early part of the sample, while during the 2008-2013 period of market distress and subsequent recovery the contribution of alternative risk premia actually exceeded that of traditional risk premia. This further highlights the crucial role that these strategies play in understanding and emulating the returns of hedge funds. In the later part of the sample, traditional risk premia outrank alternative risk premia in their contribution to hedge fund returns because of their higher degree of directionality in this long-trending market environment.

Exhibit 12: Factor Attribution of Hedge Fund Index Performance over Time

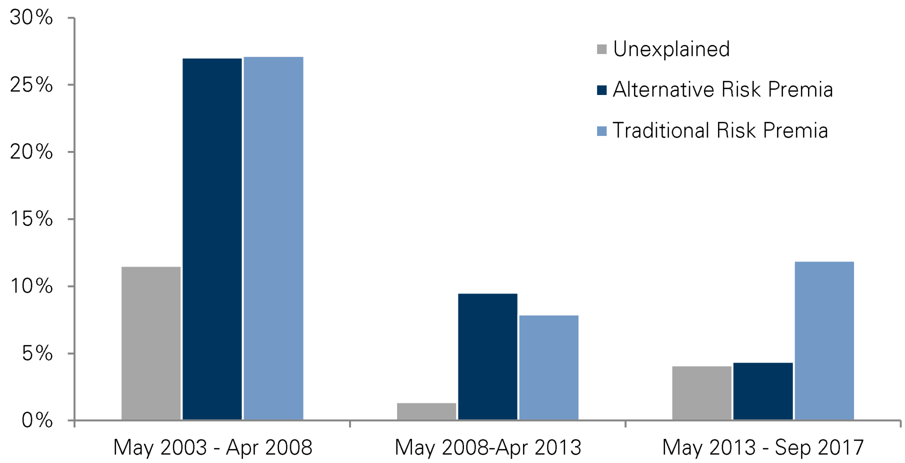

Source: HFR, BarclayHedge, GSAM as of December 2017

The final part of the analysis of returns and risk of the Hedge Fund Index applies this contribution analysis to the four main hedge fund categories over the full sample period. Focusing on the contribution of unexplained returns, Exhibit 13 indicates that the overall return impact of unexplained returns on the hedge fund index is predominantly concentrated in the Equity Long Short and Relative Value categories. Furthermore, the exhibit points to clear disparities in terms of the contribution of traditional risk premia relative to alternative risk premia across the four hedge fund categories. Whereas Equity Long Short is the most extreme case with an approximately 85/15 split of the proportional contribution in favor of traditional risk premia, the Macro category is at the other extreme with a 95/5 split of the proportional contribution in favor of alternative risk premia relative to traditional risk premia. Compared to these extremes, Relative Value's return contribution comes out very evenly between the three components.

The marginal contribution to risk by hedge fund category confirms the effect from the overall risk contribution analysis. Across all four categories, the relative role played by traditional risk premia to explain risk increases relative to the role played by alternative risk premia because of their higher inherent volatility. An additional noteworthy point relates to the relative proportion of volatility related to unexplained returns. Among the four categories, Relative Value turns out to have the highest proportional contribution, which hints at the complexities of identifying appropriately liquid vehicles to represent the complex and illiquid risk exposures hedge funds in this category tend to take.

Exhibit 13: Factor Attribution of Hedge Fund Index Performance for Individual Hedge Fund Categories (April 2003 – September 2017)

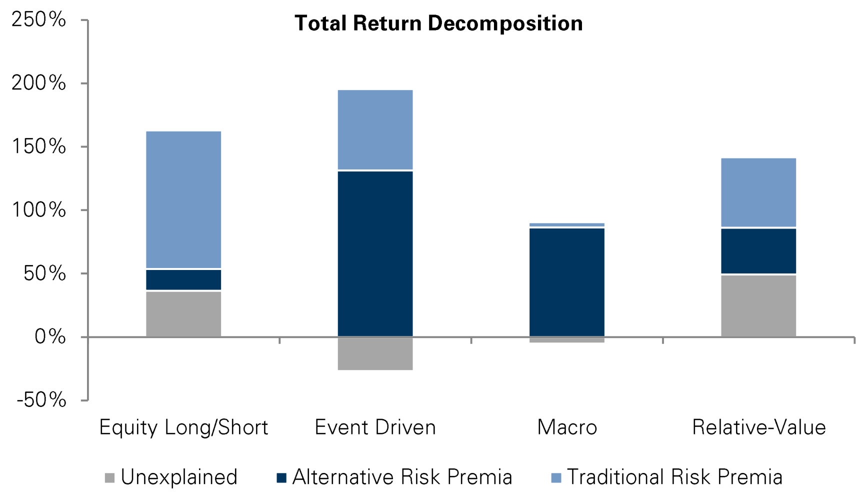

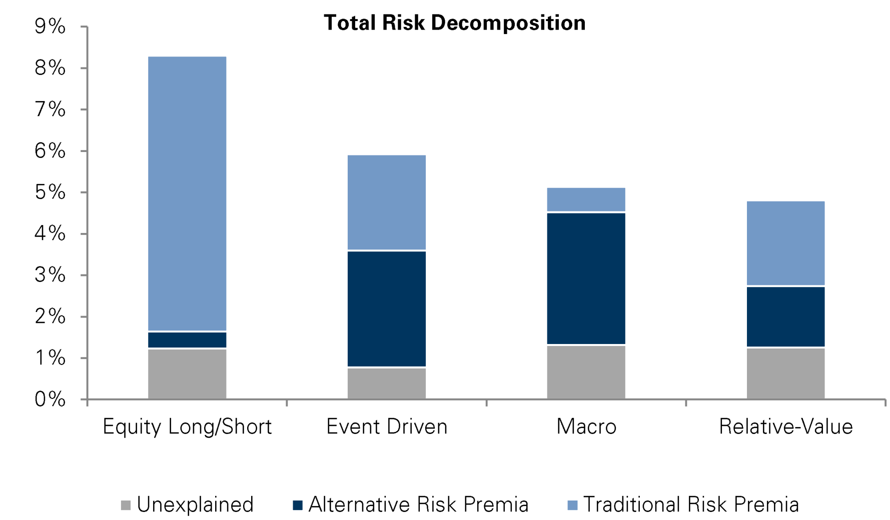

Source: HFR, BarclayHedge, GSAM as of December 2017

5 Developments in the Hedge Fund Industry

We conclude with some perspectives on the hedge fund industry, in particular their fee structure and overall liquidity. We also provide a forward-looking perspective on some near-term developments for the hedge fund industry.

5.1 The Evolution of Hedge Fund Characteristics

Fees are at the forefront of every investor's allocation decision, particularly in relation to the performance that the corresponding investment vehicle may offer and has historically realized. The question arises to what degree fee pressures may have also found their way into the hedge fund industry.

Exhibit 14: Cross-Sectional Averages of Incentive and Management Fees across Hedge Funds

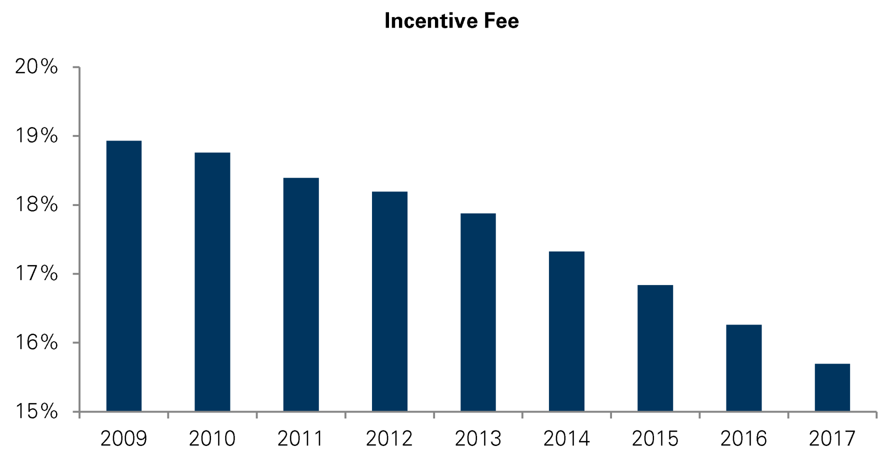

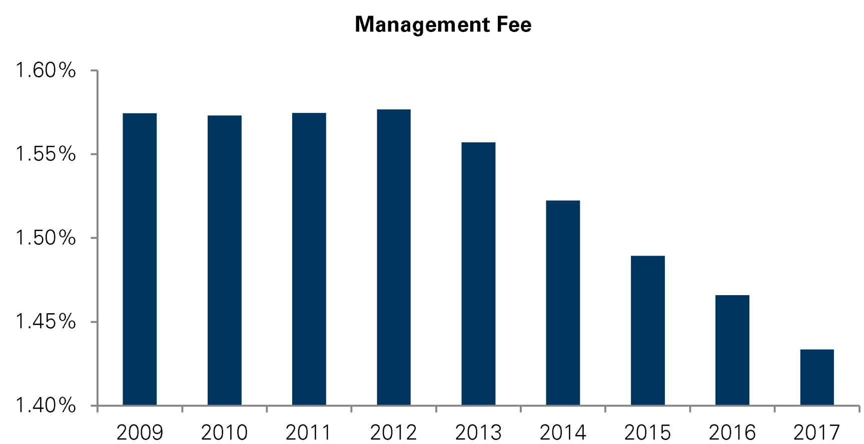

Source: HFR, BarclayHedge, GSAM as of December 2017

Exhibit 14 focuses on the fees that hedge funds charge. Their fee structure is typically composed of an incentive fee as well as a management fee. The incentive fee is charged on the profits[26] that a hedge fund generates while the management fee is charged on the total assets under management regardless of performance.

[26] It tends to be the case that incentive fees are associated with certain threshold conditions, so-called watermarks, and incentive fees only apply to profits that exceed these watermarks. The fee overview in Exhibit 14 ignores any considerations around watermarks, as the bespoke and idiosyncratic nature of watermarks presents impediments to the cross-sectional aggregation across hedge funds.

Within the overall cross-section of hedge funds, Exhibit 14 takes an average over the observed fees across all hedge funds in our sample in a given year. It is therefore not necessarily a statement about the fee evolution of individual funds, but rather an assessment of the fee evolution of the overall hedge fund universe. As far back as 2009,[27] both the incentive fee and the management fee are below the popularly quoted fee structure of "2+20", referring to a management fee of 2% paired with an incentive fee of 20%. Moreover, fees have actually turned out to be on a generally downward sloping trajectory. Incentive fees have shrunk from slightly below 19% to less than 16% over the span of eight years. Management fees initially proved more resilient at levels between 155 and 160 bps, but have since also succumbed to fee pressure to fall below 145bps.

Overall, Exhibit 14 points to the existence of fee pressure for hedge funds and the end of the commonly quoted "2+20" fee structure. That said, it is worth noting that fees are still noticeably higher than the typical fees charged for, say, Exchange Traded Products (ETPs) that provide passive exposure to a general equity market index or even ETPs that provide investors with access to specialized portions of the fixed income market, such as convertible bonds or bank loans.

Exhibit 15: Percentage of Hedge Funds with a Lock-Up Period

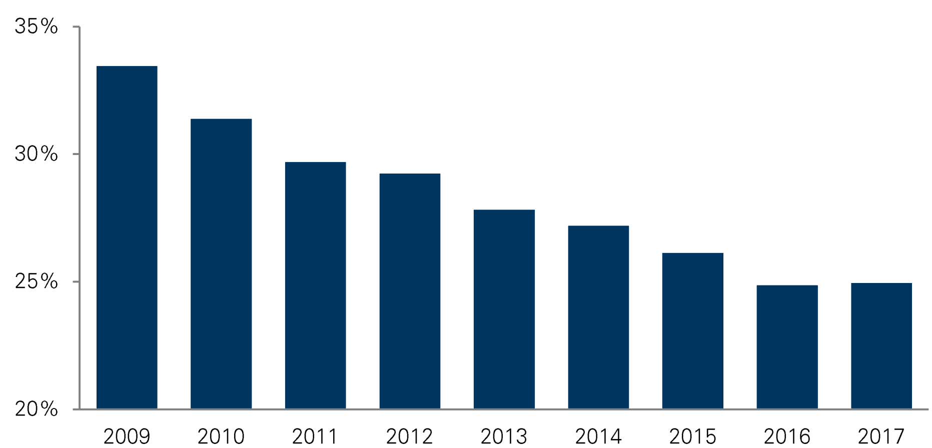

Source: HFR, BarclayHedge, GSAM as of December 2017

Another investor concern, among others, is the liquidity of their investment portfolio. In the context of hedge funds, we use the existence of a lock-up period as a proxy for liquidity. A lock-up period is typically imposed in order to enable hedge fund managers to make investments in illiquid assets and puts restrictions on the ability of hedge fund investors to redeem or sell their investments in hedge funds.

As is the case for Exhibit 14, Exhibit 15 also focuses on the overall cross-section of hedge funds and provides an assessment of the composition of the overall hedge fund universe instead of individual hedge funds. It displays the fraction of hedge funds that impose a lock-up period compared to all hedge funds in the universe that report in a given year. Over the span of eight years, the prevalence of lock-up periods has fallen continuously and now stands at 25% - almost 10% below the level in 2009, suggesting that there has been pressure on hedge funds overall to make adaptations to their liquidity restrictions.

[27] The time frame is determined by our availability of point-in-time data for the fee structure of hedge funds.

5.2 Considerations around the Implementation of Liquid Hedge Fund Tracking Strategies

Whereas section 5.1 has focused on historic, backward-looking trends in the hedge fund industry, we now aim to provide a near-term forward-looking outlook on the hedge fund universe, both in terms of performance as well as in terms of their impact in hedge fund investors' portfolios.

In terms of performance, we actually argue to move away from a narrow focus on absolute return, but advocate for a measure of risk-adjusted outperformance. Particularly given heightened fee sensitivity, hedge fund investors should at least be looking for outperformance over a fairly simplistic passive benchmark, such as a global equity market index. Because of their differing volatility levels, it is however not appropriate to compare hedge fund returns with outright returns of an equity index. Thus, we consider hedge fund returns only to the extent they outperform a beta-adjusted equity benchmark and normalize this adjusted return by the volatility of their idiosyncratic return to construct an information ratio.[28] In this sense, Exhibit 16 presents the risk-adjusted performance of the overall hedge fund universe compared to the global equity market, as represented by the MSCI World Net Total Return Index.

In line with the growth of the AUM in the overall hedge fund universe presented in Sections 1 and 2.1, hedge funds have - after adjusting for their equity beta - generated positive value over the general global equity market over the past 10+ years. However, it is also apparent that this outperformance has been far from uniform. In particular during late 2012 and 2013 and also intermittently in more recent years, hedge fund performance has been challenged, which may have led some to call into question the attractiveness of hedge funds as sources of alternative returns and has certainly had an impact on the fees that investors proved to be willing to pay and the liquidity restrictions they were willing to accept. However, the second half of 2017 has seen a sharp increase in the information ratio to levels above 1. Historically, that puts current performance into the 10th percentile of the best performing time periods going back to 2005. If this continues, questions about the attractiveness of hedge funds should decline. Given the close co-movement between the Liquid Tracking and the Hedge Fund Index, such developments also look to be potentially beneficial for the risk-adjusted returns of access vehicles to the common systematic factor exposures of the broad universe of hedge funds.

[28] Technically, we define the beta-adjusted IR as the annualized ratio of the intercept of a regression of the overall hedge fund index on the equity index and the standard deviation of the error term from this regression. Exhibit 16 displays this information ratio calculated based on a rolling 24-month window.

Exhibit 16: Beta-Adjusted IR of Overall Hedge Fund Studied Universe to MSCI World Index

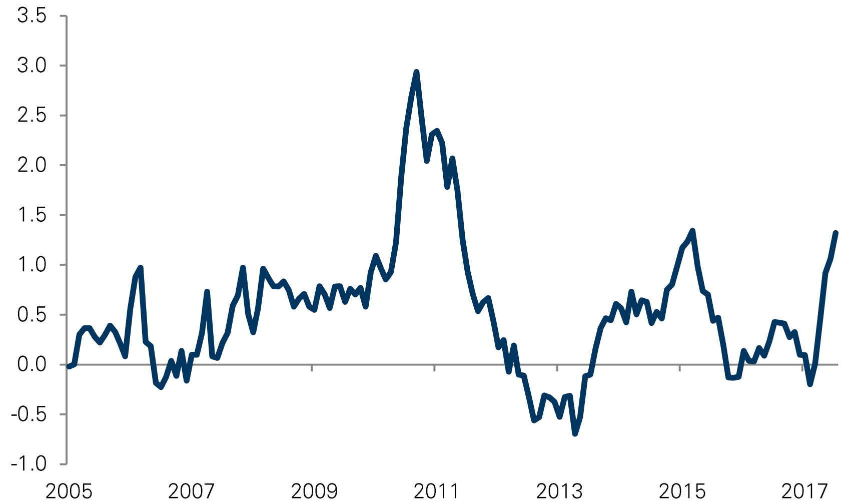

Source: MSCI, HFR, BarclayHedge, GSAM as of December 2017

Another noteworthy development in the hedge fund universe relates to diversification. Particularly towards the end of the studied sample period, diversification among hedge funds has increased dramatically to a level previously not seen in our sample that extends back to early 2003 (see what we refer to as "Diversification Ratio" in Exhibit 17). This diversification effect implies that active managers of hedge fund portfolios express more diverse views in their positions. While this effect may increase the benefits to hedge fund selection it can also increase the risks of selecting the "wrong" fund, as discussed in Section 2.2. An investor that is concerned about these types of risks might find it beneficial then to rather rely on liquid investment vehicles designed to track the returns of the hedge fund universe as a whole.

Our measure of hedge fund diversification, as displayed in Exhibit 17, is based on volatility comparisons. The volatility of the Hedge Fund Index is driven by the overall level of volatility of the hedge funds making up the universe as well as the degree to which these hedge funds are correlated to each other. It is apparent from the chart that hedge funds have generally become less volatile, as evidenced by the decline in the Average of Hedge Fund Volatilities. However, a comparison of the volatility of the Hedge Fund Index (referred to as "Volatility of Average Hedge Fund Return") to the Average of Hedge Fund Volatilities provides us with an indication of diversification between individual hedge fund returns. The more Volatility of Average Hedge Fund Return diverges from Average of Hedge Fund Volatilities, the greater the impact of diversification or lack of correlation. In this case, the Volatility of Average Hedge Fund Return has fallen more sharply than the Average of Hedge Fund Volatilities providing evidence for increased diversification.[29] The capability of the simulated Liquid Tracking Portfolio to approximate the returns of the hedge fund universe has however proven to be resilient to this increase in diversification, as evidenced by the 24-month correlation being with 96.1% in the 97th percentile when compared to history.

[29] Technically, the "Diversification Ratio" is defined as 1 minus the ratio of the difference of the volatility of the return average and a volatility measure that assumes uncorrelated hedge fund returns to the difference of a measure that assumes perfectly correlated hedge fund returns (average of individual hedge fund volatilities) and the uncorrelated measure.

Exhibit 17: Hedge Fund Volatility and Diversification

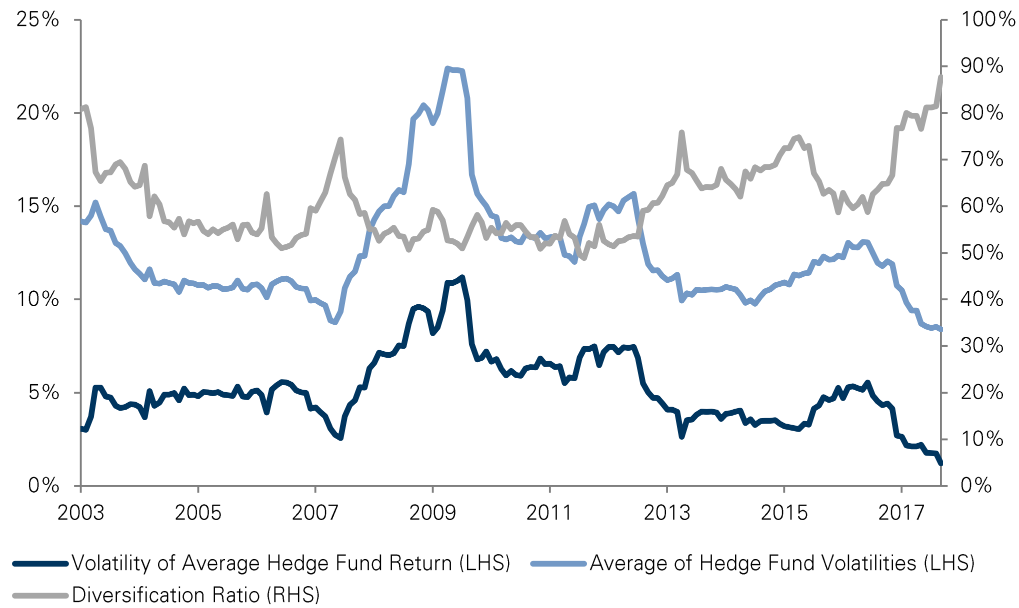

Source: HFR, BarclayHedge, GSAM as of December 2017

6 Conclusion

This article discusses an alternative to hedge fund investing based on a risk-based approach that dynamically infers the exposures to traditional and alternative risk premia present in a broad and diversified universe of hedge funds. The difference between the well-defined and liquid nature of the factors and the opaqueness and illiquidity of some hedge fund investment strategies leads to a tracking error of 2.1% between the simulated Liquid Tracking Portfolio and the aggregate performance of the hedge fund universe. However, a correlation of 93.5% between the two and the fact that 84% of hedge fund returns can be captured to an almost equal degree by exposures to traditional and alternative risk premia make this methodology a viable alternative. A potential challenge to this high degree of hedge fund return attribution in the future rests on the ongoing impact of, for example, illiquidity or non-public aspects of stock picking. Sources of hedge fund returns like these will limit the efficacy of the proposed alternative because of the reliance on liquidity and publicly available information of this approach, although historically the impact over the past 15 years has proven to be limited.

This article emphasizes the time varying nature of hedge fund positioning, as evident not only from the shifting attribution of hedge fund returns to traditional and alternative risk premia but also from the inherent dynamism of the allocations inside the alternative risk premia as well as the allocation to all risk premia. Any dynamic allocation and hedge fund positioning in particular hinges on the quality of the data to be able to monitor and assess it, which is why it is of eminent importance to have ongoing access to many and diverse sources of hedge fund information. Moreover, it is crucial to continuously enhance and refine the understanding of hedge fund investment strategies, especially through usage of alternative risk premia.

While the investment philosophy based on the identification of traditional and alternative risk premia from a broad universe of hedge fund returns is fairly unique, it generally fits into the classification of so-called “Liquid Alternative Funds” that has been created in recent years by investment research firms such as Morningstar, Inc. As of the end of 2017, there were 640 funds with aggregate AUM of $316.8 bn in this category according to an analysis based on data by Morningstar, Inc.,[30] illustrating the increased appeal of this concept to the marketplace. Over the years to come, it will be interesting to see how the interplay between hedge funds and liquid alternative funds plays out. Particularly interesting will be developments around fees, liquidity hurdles and more generally if hedge funds will be fast enough to innovate in order to generate attractive unexplained returns while an increasing amount of hedge fund know-how becomes common knowledge and finds its way into liquid alternatives funds.

7 References

Agarwal, Vikas, and Narayan Y. Naik (2000a): "Multi-Period Performance Persistence Analysis of Hedge Funds," Journal of Financial and Quantitative Analysis, Vol. 35, No. 3, pp. 327–342.

Agarwal, Vikas, and Narayan Y. Naik (2000b): "On Taking the Alternative Route: Risks, Rewards, and Performance Persistence of Hedge Funds," Journal of Alternative Investments, Vol. 2, No. 4, pp. 6–23.

Agarwal, Vikas, and Narayan Naik (2004): "Risk and Portfolio Decisions Involving Hedge Funds," Review of Financial Studies, Vol. 17.

Agarwal, Vikas, Kevin A. Mullaly, and Narayan Naik (2015): "The Economic and Finance of Hedge Funds: A Review of the Academic Literature," Foundations and Trend in Finance, Vol. 10, No. 1, pp. 1-111.

Amenc, Noel, Sina El Bied, and Lionel Martellini (2003): "Predictability in Hedge Fund Returns," Financial Analysts Journal, Vol. 59, No. 5, pp. 32–46.

Amenc, Noel, Walter Gehin, Lionel Martellini, and Jean-Christophe Meyfredi (2008): "Passive hedge fund replication: A critical assessment of existing techniques," Journal of Alternative Investments, Vol. 11, No. 2, pp. 69-83.

Amenc, Noel, Lionel Martellini, Jean-Christophe Meyfredi, and Volker Ziemann (2010): "Passive hedge fund replication – beyond the linear case," European Financial Management, Vol. 16, No. 2, pp. 191-210.

Bares, Pierre-Antoine, Rajna Gibson, and Sebastien Gyger (2003): "Performance in the Hedge Funds Industry: An Analysis of Short and Long-Term Persistence," Journal of Alternative Investments, Vol. 6, No. 3, pp. 25–41.

Bollen, Nicolas, and Gregg Fisher (2014): "Send in the clones? Hedge fund replication using futures contracts," Journal of Alternative Investments, Vol. 16, No.2, pp. 80-95.

Boyson, Nicole (2008): "Hedge fund performance persistence: A new approach," Financial Analysts Journal, Vol. 64, No.6, pp. 27-44.

[30] The aggregate of 640 funds is decomposed of 551 funds across 17 different categories that Morningstar, Inc. categorizes as liquid alternatives as well as 89 funds in the Nontraditional Bond category, which contains many funds that qualify as alternative strategies. See the Goldman Sachs Asset Management publication Liquid Alternative Investments MAPS Year End 2017.

Brown, Stephen, and William Goetzmann (2003): "Hedge funds with style," Journal of Portfolio Management, Vol. 29, No. 2, pp. 101-112.

Cai, Li, and Bing Liang (2012): “Asset allocation dynamics in the hedge fund industry,” Journal of Investment Management, Vol. 10, No. 2.

Capocci, Daniel, and Georges Huebner (2004): "Analysis of hedge fund performance," Journal of Empirical Finance, Vol. 11, No.1, pp. 55-89.

Capocci, Daniel, Albert Corhay, and Georges Huebner (2005): "Hedge fund performance and persistence in bull and bear markets," The European Journal of Finance, Vol. 11, No.5, pp. 361-392.

Carhart, Mark (1997): “On Persistence in Mutual Fund Performance,” Journal of Finance, Vol. 52, No. 1, pp. 57-82.

Edwards, Franklin, and Mustafa Caglayan (2001): "Hedge fund performance and manager skill," Journal of Futures Markets, Vol. 21, No.11, pp. 1003-1028.

Eling, Martin (2009): “Does Hedge Fund Performance Persist? Overview and New Empirical Evidence,” European Financial Management, Vol. 15, No. 2, pp. 362–401.

Fama, Eugene and Kenneth French (1992): “The Cross-Section of Expected Stock Returns,” Journal of Finance, Vol. 47, No. 2, pp. 427–465.

Fama, Eugene and Kenneth French (1993): “Common Risk Factors in the Returns on Stocks and Bonds,” Journal of Financial Economics, Vol. 33, No. 1, pp. 3–56.

Fung, William, and David Hsieh (1997): “Empirical Characteristics of Dynamic Trading Strategies,” Review of Financial Studies, Vol. 10, No. 2, pp. 275-302.

Fung, William, and David Hsieh (2001): “The Risk in Hedge Fund Strategies: Theory and Evidence from Trend Followers,” Review of Financial Studies, Vol. 14, No.2, pp. 313-341.

Fung, William and David A. Hsieh (2004): “Hedge Fund Benchmarks: A Risk-Based Approach,” Financial Analysts Journal, Vol. 60, No. 5, pp.65-80.

Hill, Joanne M., Barbara Mueller, and Venkatesh Balasubramanian (2004): “The ‘Secret Sauce’ of Hedge Fund Investing – Trading Risk Dynamically.” Goldman Sachs Equity Derivatives Strategy, pp. 1-24.

Hasanhodizic and Lo (2007): “Can Hedge-Fund Returns Be Replicated?: The Linear Case,” Journal of Investment Management, 5(2), pp. 5-45.

Jaeger, Lars (2008): “Alternative Beta Strategies and Hedge Fund Replication,” John Wiley & Sons, Hoboken.

Jegadeesh, Narasimhan, and Sheridan Titman (1993): “Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency,” Journal of Finance, Vol. 48, No. 1, pp. 65-91.

Joenvaara, Juha, Robert Kosowski and Pekka Tolonen (2016): “Hedge Fund Performance: What Do We Know?”, Working Paper, SSRN.

Kat, Harry, and Helder Palaro (2005): "Who Needs Hedge Funds? A Copula-Based Approach to Hedge Fund Replication," Working Paper.

Liang, Bing (1999): "On the performance of hedge funds," Financial Analysts Journal, Vol. 55, No.4, pp. 72-85.

Lintner, John (1965): "The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets," Review of Economics and Statistics, 47 (1), pp. 13–37.

Malkiel, Burton, and Atanu Saha (2005): "Hedge Funds: Risk and Return," Financial Analysts Journal, Vol. 61, No. 6, pp. 80–88.

Mossin, Jan (1966): "Equilibrium in a Capital Asset Market," Econometrica, Vol. 34, No. 4, pp. 768–783.

Patton, Andrew, and Tarun Ramadorai (2013): "On the High-Frequency Dynamics of Hedge Fund Risk Exposures," Journal of Finance, Vol. 68, No. 2, pp. 597-635.

Sharpe, William F. (1964): "Capital asset prices: A theory of market equilibrium under conditions of risk," Journal of Finance, 19 (3), pp. 425–442.

Sharpe, William F. (1992): "Asset Allocation: Management Style and Performance Measurement," Journal of Portfolio Management, Vol. 18, No. 2, pp.7-19.

Schneeweis, Thomas, and Richard Spurgin (1998): "Multifactor analysis of hedge fund, managed futures, and mutual fund return and risk characteristics," Journal of Alternative Investments Vol. 1, No. 2, pp. 1-24.

Scholes, Myron, and Joseph Williams (1977): "Estimating betas from nonsynchronous data," Journal of Financial Economics, Vol. 5, pp. 309-327.

Ter Horst, Jenke, and Marno Verbeek (2007): "Fund liquidation, self-selection, and look-ahead bias in the hedge fund industry," Review of Finance, Vol. 11, No.4, pp. 605-632.

General Disclosures

This material is provided at your request for informational purposes only. It is not an offer or solicitation to buy or sell any securities.

This material is provided for educational purposes only and should not be construed as investment advice or an offer or solicitation to buy or sell securities.

THIS MATERIAL DOES NOT CONSTITUTE AN OFFER OR SOLICITATION IN ANY JURISDICTION WHERE OR TO ANY PERSON TO WHOM IT WOULD BE UNAUTHORIZED OR UNLAWFUL TO DO SO.

These examples are for illustrative purposes only and are not actual results. If any assumptions used do not prove to be true, results may vary substantially.

Backtested performance shown is not actual performance and in no way should be construed as indicative of future results. Backtested performance results are created based on an analysis of past market data with the benefit of hindsight, do not reflect the performance of any GSAM product and are being shown for informational purposes only. Please see additional disclosures.

Simulated Performance

Simulated performance is hypothetical and may not take into account material economic and market factors, such as liquidity constraints, that would impact the adviser's actual decision-making. Simulated results are achieved by retroactively applying a model with the benefit of hindsight. The results reflect the reinvestment of dividends and other earnings, but do not reflect fees, transaction costs, and other expenses a client would have to pay, which would reduce returns. Actual results will vary.

Index Benchmarks

Indices are unmanaged. The figures for the index reflect the reinvestment of all income or dividends, as applicable, but do not reflect the deduction of any fees or expenses which would reduce returns. Investors cannot invest directly in indices.

The indices referenced herein have been selected because they are well known, easily recognized by investors, and reflect those indices that the Investment Manager believes, in part based on industry practice, provide a suitable benchmark against which to evaluate the investment or broader market described herein. The exclusion of "failed" or closed hedge funds may mean that each index overstates the performance of hedge funds generally.

The website links provided are for your convenience only and are not an endorsement or recommendation by GSAM of any of these websites or the products or services offered. GSAM is not responsible for the accuracy and validity of the content of these websites.

This material is provided for educational purposes only and should not be construed as investment advice or an offer or solicitation to buy or sell securities.

This information discusses general market activity, industry or sector trends, or other broad-based economic, market or political conditions and should not be construed as research or investment advice. This material has been prepared by GSAM and is not financial research nor a product of Goldman Sachs Global Investment Research (GIR). It was not prepared in compliance with applicable provisions of law designed to promote the independence of financial analysis and is not subject to a prohibition on trading following the distribution of financial research. The views and opinions expressed may differ from those of Goldman Sachs Global Investment Research or other departments or divisions of Goldman Sachs and its affiliates. Investors are urged to consult with their financial advisors before buying or selling any securities. This information may not be current and GSAM has no obligation to provide any updates or changes.

Although certain information has been obtained from sources believed to be reliable, we do not guarantee its accuracy, completeness or fairness. We have relied upon and assumed without independent verification, the accuracy and completeness of all information available from public sources.