the journal of
PORTFOLIO management

PORTFOLIO MANAGEMENT RESEARCH
with. Intelligence

August 2025 | Volume 51 | Number 9

pm-research.com

Strategic Asset Allocation with Alternative Investments:
An Integrated Approach

Alexander Rudin and Daniel Farley

STATE STREET GLOBAL ADVISORS

Alex Rudin, Ph.D.

Senior Managing Director, Head of Investment Research – Investment Solutions Group and Fixed Income

Alex leads quantitative research efforts across State Street Investment Management's Multi-Asset ("Investment Solutions Group") and Fixed Income Cash and Currency ("FICC") business lines. Within ISG, Alex drives the quantitative research agenda in support of tactical and strategic asset allocation processes, portfolio construction analytics, target volatility models, and the target date fund franchise. Within FICC, Alex oversees quantitative research efforts for enhanced indexing, systematic credit, ETF management, and top-down active fixed income process.

In addition, Alex works with the broader ISG and FICC business teams to help bring new ideas, products, and thought leadership to our clients.

Alex joined State Street Investment Management in September 2014. Before joining ISG in 2018 to head up its research effort, Alex was Global Head of Liquid Alternative Investments at State Street Investment Management.

Alex holds a Ph.D. in Theoretical Physics, has 24 years of industry experience in quantitative finance and alternative investments, and is an author or co-author of more than 40 articles in academic journals in the areas of theoretical physics and finance.

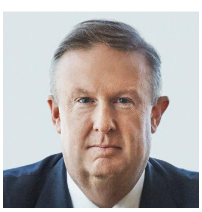

Daniel Farley

Dan is an Executive Vice President of State Street Investment Management and CIO of the Investment Solutions Group. In this role, he oversees a global team of over 75 investment professionals managing over US$266B in multi asset class portfolios. His team is responsible for the firm's target date, active asset allocation, and Outsourced CIO investment strategies. He is also a member of the firm's Executive Management Group. Dan joined the firm in 1992 and has over 25 years of investment experience. Prior to this role, he was responsible for the US multi asset class solutions team.

Dan holds an MBA from Bentley University, a BSBA from Stonehill College and has earned the Chartered Financial Analyst (CFA) designation. He is a member of the CFA Institute and the CFA Society Boston. He is executive sponsor of the firm's Latin American Professionals Group. Dan is a frequent speaker with the media and conferences on a variety of investment topics. He is also on the Board of Trustees at Bentley University and the Chairman of the Board at the Crispus Attucks Children's Center.

Strategic Asset Allocation with Alternative Investments: An Integrated Approach

Alexander Rudin and Daniel Farley

Alexander Rudin
is senior managing director, head of Investment Research – Investment Solutions Group and Fixed Income at State Street Investment Management in Boston, MA.
alexander_rudin@ssga.com

Daniel Farley
is the chief investment officer of Investment Solutions at State Street Investment Management in Boston, MA.
daniel_farley@ssga.com

KEY FINDINGS

- In the vast majority of cases, performance of alternative investments is closely linked to that of their public counterparts; recognizing this connection is a necessary step toward self-consistent strategic asset allocation.
- It is much easier to forecast risk properties of idiosyncratic investments such as hedge funds and private equity vehicles than their future returns; a version of Bayesian shrinkage applied to returns could be used to recognize this empirical fact.
- Historical evidence points to private assets such as private equity and credit being well-modeled as leveraged versions of their public counterparts, but without the short-term excess volatility component embedded in public prices. A long-horizon modification of the traditional strategic asset allocation framework enables a consistent treatment of this phenomenon and harmonizes return/risk properties across long-only and alternative assets.

ABSTRACT

Quantitative techniques for incorporating alternative investments and particularly private assets into strategic asset allocation remain unsettled and insufficiently covered in the literature. This article attempts to bridge this gap by offering a broad framework for such portfolio construction tasks. The framework incorporates a range of techniques the authors see as crucial, including alpha/beta separation, Bayesian shrinkage, return unsmoothing, long-horizon risk estimation, and event-risk constraints. The author's suggested approach places all types of assets, ranging from long-only to hedge funds to private equity and credit, on an even playing field by recognizing commonalities in their return drivers while also treating the idiosyncratic elements of those same returns in a way that is appropriate for each asset class.

Several books have been written about quantitative aspects of hedge fund investing, with multiple chapters in each dedicated to portfolio construction (see, for example, Lhabitant 2004 and Molyboga and Swedroe 2023). The legitimate questions are, what new can be added to this well-studied matter, and can a single relatively short article add value to extensive volumes already written? We hope that the answer is yes to both questions.

The purpose of this article is to provide a practical guide on how to construct portfolios that include alternative investments. Instead of reviewing an exhaustive list of available techniques, we will start by pointing out what makes portfolio construction with alternatives so challenging and how those challenges narrow the number of quantitative tools one can use effectively. We will also discuss how an allocator can think of alternatives in a broad portfolio and translate these considerations into an actionable asset allocation process across traditional long-only assets, hedge funds, private equity, and private credit. While reviewing several approaches popular among practitioners, we will take an in-depth dive into an approach we call *integrated strategic asset allocation*, or *integrated SAA*. Despite being straightforward to implement, the integrated SAA methodology reflects fundamental properties of these asset classes and how they interact with the traditional ones. This, in our view, makes integrated SAA highly appropriate for asset allocation practitioners.

The rest of the article is structured as follows. First, we delve into the theoretical aspects of hedge fund return decomposition. Next, we briefly review various approaches practitioners use for asset allocation with alternatives, followed by an in-depth explanation into integrated SAA. We then address another challenge associated with alternative investments—and particularly with private assets—and that is the smoothness of their returns. We discuss several unsmoothing techniques, including those that can simultaneously unsmooth and determine factor exposures. After covering unsmoothing, we discuss differences and commonalities between public and private asset returns, emphasize importance of time horizon for risk estimation, and introduce methodological ways to make it consistent across asset classes. Then we describe dual-horizon integrated SAA as a way to incorporate long-only assets, hedge funds, and private equity and credit into the same portfolio in a self-consistent, intellectually coherent fashion. Finally, we summarize and conclude.

SEPARATING ALPHA FROM BETA IN HEDGE FUND RETURNS

The term *hedge funds* covers an extremely broad range of investment strategies. These strategies vary widely in terms of underlying instruments and drivers of return. On top of that, hedge funds often vary their management style over time and are usually opaque in terms of holdings. In other words, they represent a challenge for a quantitatively inclined asset allocator who is used to working with Markowitz-style portfolio construction philosophy with its ingrained dependency on well-defined expectations for returns, risks, and correlations for all assets.

Fortunately, the objective of a practical investor is not to find an absolutely correct solution, but the most reasonable one given the information available. So, what do we know about *typical*[1] hedge fund investments? We point out two broad observations.

First is that although there is massive differentiation between hedge fund styles in terms of return drivers, *within* a given style, managers often have a substantial amount of commonality (see, for example, Pedersen, Page, and He 2014, Rudin 2018, and the literature review therein). Most critical of those commonalities is some level of exposure to known risk premiums, such as public equities or credit, or to something more exotic but still well understood, such as trend following or defensive options overlays.

The second observation is that the amount of historical performance data available for a typical hedge fund is very small by the conventional data science standards (most hedge fund have less than 10 years of history, which corresponds to fewer than 120 monthly data points). This severely limits our ability to glean information from hedge fund historical return streams and leaves most of a sophisticated quantitative investor's analytical tools out of practical use.

[1] We emphasize the word *typical*, as there will always be hedge funds that do not fit neatly into the provided description.

EXHIBIT 1

Schematic Attribution of Hedge Fund Returns

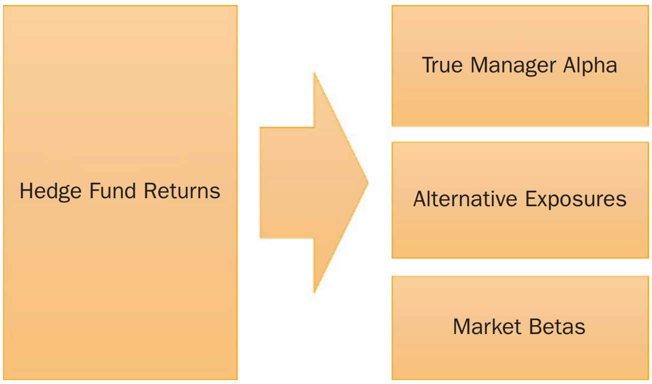

Let's now consider practical implications of these observations. At a high, simplistic level, one can attribute hedge fund returns to three broad sources (see Exhibit 1). The first one is long-only equity, rates, or credit exposure—we name that source *market betas*. Second is some sort of a dynamic, but well-understood, transparent, and easily replicable investment strategy such as trend following or options overlay. We name this second category *alternative exposures*. The remainder of the return streams of managers is, of course, unique, and we name it *true alpha*.

From the quantitative portfolio construction perspective, being unique is not an attractive feature. Any portfolio construction process must necessarily rely on some explicit or implicit assumptions about the future behavior of assets. That brings us to a key question: Which attributes of historical return streams of hedge fund assets can be relied upon for portfolio construction purposes and, conversely, which attributes are not persistent and hence need to be de-emphasized or replaced by a group average?

Rudin (2018) did an interesting study of this matter. The author attempted to separate returns of broad groups of hedge funds (following the same investment style) into the systematic component and the idiosyncratic component and then understand the persistence of both.

Methodologically, the process proceeded as follows. We assume that we have a number of hedge funds with a historical track record of the *i*th fund provided in a form of the time series $r_t^{(i)}$. Denoting *common factor* history as F_t, one can attempt to express hedge fund returns as

$$r_t^{(i)} = \beta_i F_t + \alpha^{(i)} + \varepsilon_t^{(i)} \tag{1}$$

Here, β_i describes exposure of the *i*th hedge fund to that common factor, and $\alpha^{(i)} + \varepsilon_t^{(i)}$ yields the idiosyncratic, residual alpha part of the return. $\alpha^{(i)}$ is a constant, whereas $\varepsilon_t^{(i)}$ is a residual return stream with a mean of zero and a volatility of σ_i.

Subsequent analysis in Rudin (2018) focused on the persistence of three key features of hedge fund returns within a given style:

- Persistence of fund exposure to the common factor (the beta)
- Persistence of idiosyncratic risk exposure (volatility of residual returns)
- Persistence of manager alpha itself, on a risk-adjusted basis[2]

Omitting details that can be found in Rudin (2018), the experiment progressed as follows. Following some basic data availability requirements, we selected a large group of managers within a hedge fund style. Then we analyzed each individual manager's track record on a moving window basis. We chose the *study* window to have 48 months' worth of historical monthly returns, and we chose the *forecast* window as a 24-month period immediately following the study window. From the historical data, we calculated and recorded market betas and residual volatilities and information ratios for the study and forecast windows. Then, we shifted moving windows by a certain time interval and repeated the analysis. For each individual manager, the process stopped when the end of the historical sample was reached.

[2] Specifically, that of $I_i = \alpha_i / \sigma_i$, minus the average of that ratio for the peer group.

EXHIBIT 2

Results for Two Popular Hedge Fund Styles: Equity Long/Short and Global Macro/Managed Futures

Panel A: Empirical Study of Persistence of Common Factor Betas, Excess Return Volatilities, and Manager Risk-Adjusted Outperformance for Equity Long/Short Investment Style

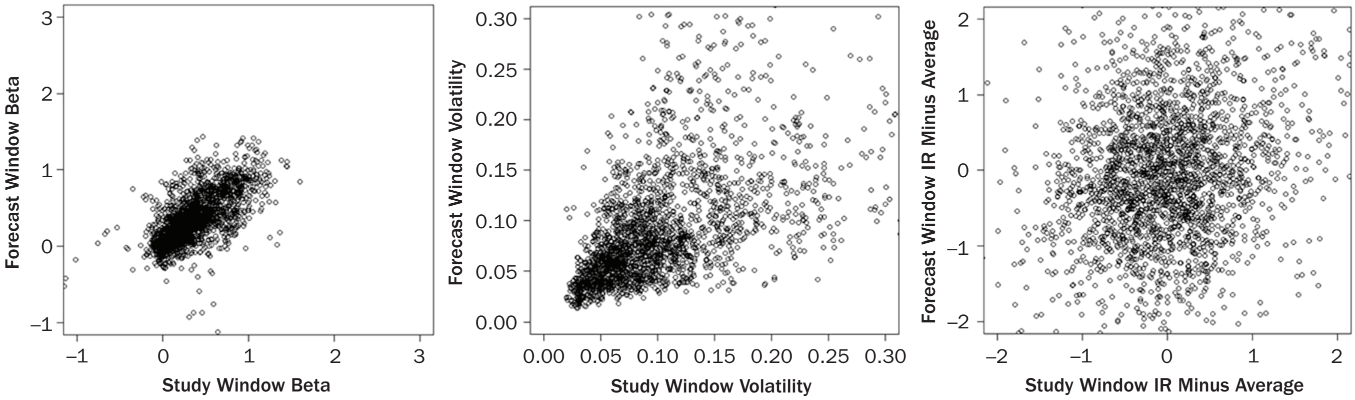

Panel B: Empirical Study of Persistence of Common Factor Betas, Excess Return Volatilities, and Manager Risk-Adjusted Outperformance for Global Macro/Managed Futures Investment Style

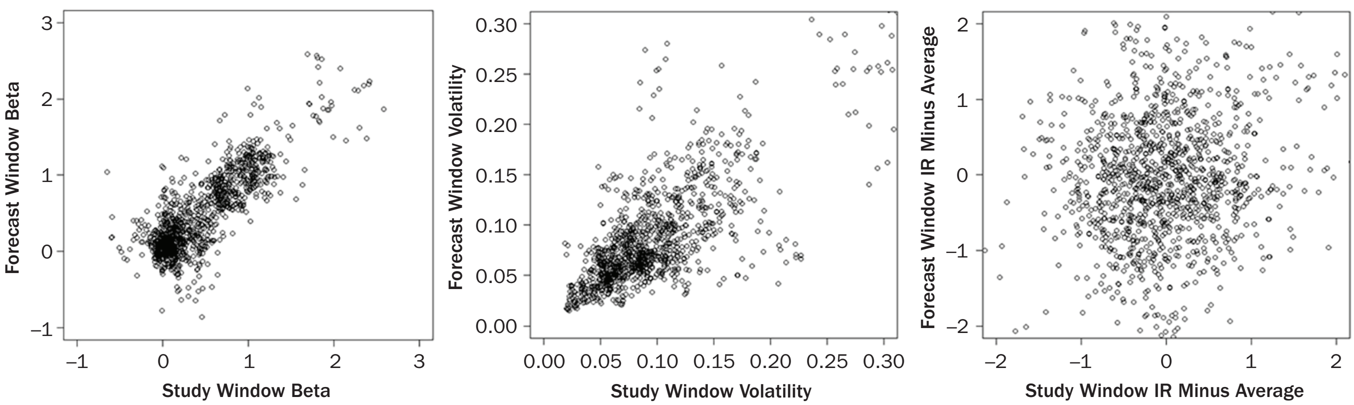

NOTE: We defined the common factor for Panel A as the Russell 2000 Index, and the SG Trend Following Index for Panel B.

Reviewing Exhibit 2, Panels A and B, strongly supports the observations we previously outlined. Within a given investment style, hedge funds shared *known*, common exposures—betas are clearly biased toward positive values. The amount of such common exposure for most managers (the beta) is highly persistent over time and thus can be successfully gleaned from historical data. Hedge funds also manage idiosyncratic risks quite well: The volatility of hedge fund residual returns (returns stripped of common factor exposures) is also persistent.

That said, there is no visible persistence of excess return information ratios in manager returns when one subtracts the grand average component, as we did. *Past performance is not at all indicative of future results.*[3]

We covered returns and volatilities, and hedge fund residual returns, but what about correlations between those residual returns? Rudin (2018) studied this aspect as well. The results were as highly consistent as they could be disappointing: On average,

[3] It is worth noting that these findings do not at all imply that manager selection is unhelpful in portfolio management. Manager selection skill is crucial to finding hedge funds that, as a group have a higher (gross) information ratio of their strategies than the industry as whole while delivering those strategies at competitive fees. What we are pointing out is that over a realistic institutional investment horizon (two to five years), tactical management of hedge fund portfolios based on short-term out- or underperformance of managers vis-à-vis their peers is extremely challenging and usually adds no value.

historically estimated correlation matrixes for residual returns were indistinguishable from random matrixes of the same dimension. That makes their historical estimates unreliable.

These empirical findings fit nicely with a massive body of academic and practitioner's research (see, for example, Brandt 2010 for an extensive earlier literature review, and Bun 2016, Rudin 2018, and Rudin, Mor, and Farley 2020 for some recent research advances), indicating that deriving correlations and information ratios directly from historical sample estimates leads to *error maximization* and portfolios that are grossly suboptimal out of sample. Instead, various forms of risk budgeting, Bayesian *shrinkage* techniques, and model-based (as opposed to history-based) future return estimation are recommended for portfolio construction. How to create a self-consistent strategic asset allocation framework that includes hedge fund assets with all their peculiarities is the topic we are turning to next.

SAA WITH HEDGE FUNDS—AN INTRODUCTION

For almost all institutions, some sort of strategic asset allocation (SAA) framework sits at the core of the investment process. With full appreciation that we are generalizing and simplifying, such SAA framework usually consists of three components: (1) capital market assumptions (CMA) for asset class returns, (2) some sort of multi-asset risk model that aims to measure risks and forecasting errors, and (3) an optimization engine—often mean–variance based—that puts all of this together in a self-consistent fashion. Exhibit 3 visualizes such structure.

Whenever an institutional investor decides to add hedge fund exposure, he faces an immediate challenge. It is difficult to develop expected returns for those complicated and idiosyncratic assets that are needed to fit them into the traditional SAA. Exhibit 4 visualizes some common approaches that allocators use to meet that challenge.

Some investors sidestep the challenge by treating hedge funds as an *add-on allocation* that sits next to the existing SAA (for example, they allocate 10% of the total portfolio to hedge funds, and the rest of the SAA remains proportionally unchanged). We argue that this approach is crude and suboptimal; many hedge funds contain long-biased exposures to equities, rates, and credit, which when taken into consideration could either impact the types of hedge funds selected relative to the remainder of the portfolio or impact allocation to non-hedge fund assets in a more optimal way. To the extent the investor has views on those asset classes expressed through the capital market assumptions, add-on allocations could be inconsistent with those views.

EXHIBIT 3

A Typical Strategic Asset Allocation Process

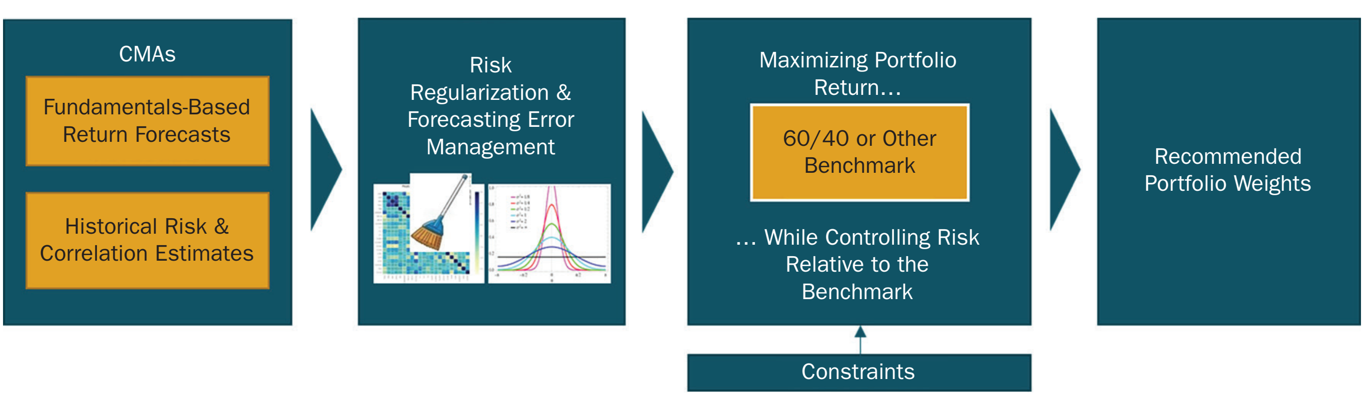

EXHIBIT 4

Role of Hedge Funds in SAA—A Comparison of Approaches

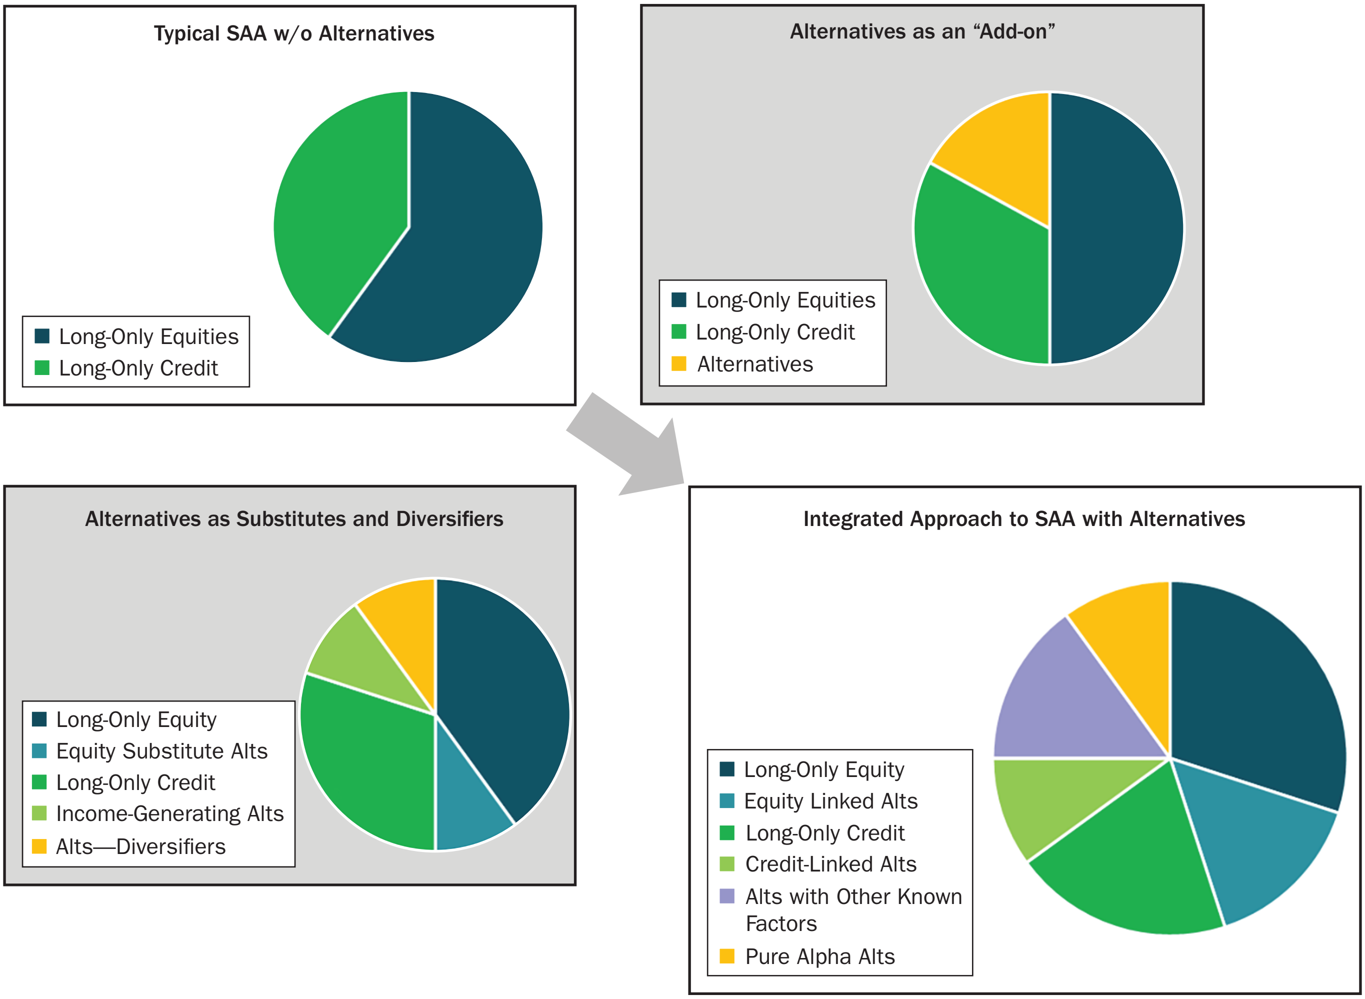

Other investors use hedge funds as the asset class *substitutes* and *diversifiers*. This is a step in the right direction, as it implies hedge fund–style differentiation by return drivers.

For example, what falls into the *equity substitute* category is implied to be correlated to long-only equity in terms of return properties. This is generally correct, but in practice that correlation may vary from negative (for short-sellers) to neutral (for many quantitative equity managers) to moderately positive (majority of fundamental long–short equity) to highly positive (activists and leveraged long-biased). In the end, a portfolio made of substitute managers often has its aggregate equity beta at a level that is far lower than long-only equity investments, leading to underperformance during periods of strong equity appreciation.

The situation with *credit substitute* is even more difficult, as managers with both *credit* focus and *income* focus may come into play. The latter category is quite broad, and alongside genuine credit involves dividend equities, fixed income relative value, and some other, highly differentiated hedge fund styles with varying degrees of exposures to rates, equities, and spreads. The same applies to the generic diversifier category.

Although the substitutes and diversifiers approach is a step in the right direction, its drawbacks are evident. Overall, SAA exposure to key return drivers such as equities, rates, and credit spreads remains fuzzy and poorly tracked, potentially leading to portfolios fundamentally inconsistent with capital market assumptions and risk levels the investor wishes to implement.

The approach we prefer enables more precise accounting of various exposures. We call it integrated, as it aims to combine traditional and alternative assets into a single, self-consistent framework.

SAA WITH HEDGE FUNDS—AN INTEGRATED APPROACH

Let us now design a variation of a traditional portfolio construction process that is consistent with both the traditional SAA framework (Exhibit 3) and the empirical facts about hedge fund returns that we described earlier.

Analytically, most SAA frameworks schematically pictured in Exhibit 3 rely on a version of a mean–variance framework with a classical quadratic utility function:

$$U(\boldsymbol{w}) = \mu^T \boldsymbol{w} - \lambda \frac{1}{2} \boldsymbol{w}^T \Sigma \boldsymbol{w} \tag{2}$$

Here $\boldsymbol{w}$ is a vector of portfolio weights, μ is a vector of expected returns, Σ is the covariance matrix, and λ is a risk aversion coefficient. For a benchmark-relative formulation, the expression remains unchanged, but $\boldsymbol{w}$, μ, and Σ correspond to active weights, relative returns, and relative return covariances, respectively.

For the traditional long-only assets, expected returns are usually provided by a proprietary capital markets assumption model, whereas the covariance matrix is constructed from historical data, usually with an element of *cleaning* (see, for example, Bun 2016 for methodology options).

If we choose a model for hedge fund returns along the lines of Equation 1, adding those complicated assets to the optimization framework becomes straightforward if we make the following assumptions:

1) We assume that we are able to identify factors (F_t^K) that are common to most managers in the given hedge fund investment style, form capital market assumptions for those factors ($\widehat{F^K}$), and calculate exposure of individual managers to those factors (β_i^K).
2) We assume that the expected information ratio (that of residuals, or *excess returns*) for all managers within a particular group or style is the same[4] and equal to the average of the information ratios across all those managers $\left(\bar{I} = \frac{1}{N} \sum_i \frac{\alpha^{(i)}}{\sigma_i}\right)$.
3) We assume that correlations across residuals are indistinguishable from zero, and the covariance matrix for residuals can be shrunk to the diagonal form ($\Omega_{ij} = \left\langle \varepsilon_t^{(i)} \varepsilon_t^{(j)} \right\rangle = \sigma_i^2 \delta_{ij}$. We notate $\delta_{ij} = 1$ if $i = j$ and $\delta_{ij} = 0$ otherwise.

These assumptions allow us to build SAA inputs for hedge fund assets in a way fully consistent with the traditional factor exposures, to the extent those hedge fund assets are driven by those exposures, specifically expected returns and variances/covariances.

---

[4] Assuming that all managers have the same information ratio of residuals is in our opinion an important base case, but of course this assumption can be relaxed by introducing active views. It can also be augmented by differentiating managers by the fees they charge. See Rudin (2018).

Expected returns mean we could take our common factor expectations, beta adjust them, and then add a generic, hedge fund–style-based alpha: $\hat{r}_i = \sum_K \beta_i^K \widehat{F^K} + \bar{I}\sigma_i$. Here we would be taking advantage of the fact that most common factors are variations of traditional long-only assets such as an equity or a bond market index, for which CMAs are commonly articulated.

Variances/covariances are a mix of common factor–related volatilities and idiosyncratic components, specifically:

- For inter-hedge fund variances/covariances: $< r_i r_j >= \sum_{KM} \beta_i^K \beta_j^M X_{KM} + \sigma_i^2 \delta_{ij}$
- For covariances between hedge funds and other assets: $\langle r_i r_A \rangle = \sum_K \beta_i^K \langle F^K r_A \rangle$

Here $X_{KM} = \langle F^K F^M \rangle$ is the inter-factor covariance.

It is worth noting that this approach strongly rhymes with the concept of Bayesian shrinkage (see, for example, Brandt 2010), but when such shrinkage is applied to hedge fund residuals as opposed to hedge fund returns as a whole.

To get a better feel for this portfolio construction process, let us go through a simple example.

CASE STUDY: ADDING A BASKET OF EQUITY LONG–SHORT FUNDS TO A LONG-ONLY PORTFOLIO

Let us attempt to build an SAA that blends five core assets often found in strategic portfolios: US large-cap equities, US small-cap equities, US aggregate bonds, US high yield, and cash.

For the purposes of this case study, we chose a simple, *relative* risk mean–variance-based portfolio construction process that aims to maximize portfolio efficiency (ratio of excess returns over the benchmark to the tracking error to that benchmark) for a given level of tracking error under standard assumptions of full investment ($\sum_i w_i = 1$) and no shorting ($w_i > 0$).

We chose a 60/40 mix of US large-cap equities and US aggregate bonds as a benchmark, and 200 bps as a tracking error target.

Results are shown in Exhibit 5, Panel A, alongside our hypothetical capital market assumptions. Large-cap equity allocation was diversified by small-cap equities. High-yield asset was allocated capital from both equities and bonds, just as one would expect. Finally, some money was parked in cash. Overall efficiency of the portfolio is positive at 0.15, meaning that the changes were consistent with our declared goal of finding a portfolio that outperforms the benchmark in a way that maximizes the relative risk-adjusted performance.

Now let us bring in a basket of hypothetical hedge fund assets that we picked from a standard hedge fund database. We chose a pool of 10 hedge funds, each of which having 10 years of monthly historical performance. Our specific choices were rather arbitrary; the goal was to illustrate the concept and the process and not make any sort of specific recommendations on the portfolio composition. For illustrative purposes, we made sure that our assets exhibited a broad range of equity market exposure.

Exhibit 5, Panel B, describes observed raw (historical) return characteristics for our manager pool alongside attribution information obtained through the lens of Equation 1. We used a US small-cap index as a common factor for the analysis, as many equity long–short managers in the US invest in small to midcap spaces. The beta to this factor ranged from 0 to 0.7 in our chosen manager pool. The average information ratio of residuals for this group of managers can easily be calculated as 0.57, and we will use that number as a starting point for forming alpha assumptions across the whole manager pool. The total expected returns shown in Exhibit 5, Panel B, are a combination of equity market expected contribution (beta-adjusted) and manager alpha.

We now are in a position to repeat our SAA calculation, but with hedge funds included. Results are in Exhibit 5, Panel C (a table representation) and Panel D (a visual representation).

Comparing recommended allocations with and without hedge funds, one can notice that the initial US small-cap allocation disappeared and the US high-yield allocation reduced to make space for hedge funds. This is because, given the objectives, hedge funds presented a more efficient use of capital than small-cap stocks and

EXHIBIT 5

SAA With and Without Hedge Funds—A Hypothetical Example

Panel A: SAA Example Without Hedge Funds

Asset Class	Expected Returns	Expected Risk	SAA Without HF	Benchmark SAA	Active Weight
US Large Cap	7.0%	15.7%	42.3%	60.0%	–17.7%
US Small Cap	8.0%	20.6%	10.2%	0.0%	10.2%
US Investment Grade Bond	4.0%	5.0%	23.5%	40.0%	–16.5%
US High Yield Bond	6.0%	9.6%	22.3%	0.0%	22.3%
US Cash	2.0%	0.6%	1.6%	0.0%	1.6%
Total			100.0%	100.0%	0.0%
Expected Return			**6.1%**	**5.8%**	**0.3%**
Expected Risk			**10.5%**	**9.7%**	**2.0%**
Efficiency Ratio					**0.15**

Panel B: Return and Risk Characteristics of Our Hypothetical Hedge Fund Managers

	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	R2000
Historical Returns	6.27%	3.03%	3.99%	6.57%	12.29%	8.10%	11.22%	10.42%	8.53%	12.14%	9.87%
Historical Volatility	8.54%	10.86%	11.46%	8.91%	8.69%	10.10%	11.23%	13.19%	14.46%	15.33%	19.38%
Beta to Russell 2000	0.00	0.21	0.27	0.29	0.27	0.42	0.44	0.61	0.69	0.71	1.00
Alpha	6.27%	0.74%	1.00%	3.35%	9.08%	3.45%	6.22%	3.78%	1.29%	4.52%	
Vol of Residuals	8.54%	10.10%	10.20%	6.99%	6.95%	6.02%	7.30%	5.76%	5.42%	6.86%	
Alpha IR	0.73	0.07	0.10	0.48	1.31	0.57	0.85	0.66	0.24	0.66	
Average Alpha IR	**0.57**										
Total Expected Returns	**4.84%**	**7.37%**	**7.94%**	**6.24%**	**6.09%**	**6.76%**	**7.66%**	**8.16%**	**8.61%**	**9.55%**	**8.00%**

Panel C: Integrated SAA With Hedge Funds

Asset Classes	Expected Return	Expected Risk	SAA With HF	SAA Without HF	Benchmark SAA	Active Weight
US Large Cap	7.0%	15.7%	44.8%	42.3%	60.0%	–15.2%
US Small Cap	8.0%	20.6%	0.0%	10.2%	0.0%	0.0%
US Investment Grade Bond	4.0%	5.0%	18.4%	23.5%	40.0%	–21.6%
US High Yield Bond	6.0%	9.6%	4.0%	22.3%	0.0%	4.0%
US Cash	2.0%	0.6%	0.0%	1.6%	0.0%	0.0%
M1	4.8%	8.5%	2.7%		0.0%	2.7%
M2	7.4%	10.9%	3.0%		0.0%	3.0%
M3	7.9%	11.5%	3.9%		0.0%	3.9%
M4	6.2%	8.9%	1.7%		0.0%	1.7%
M5	6.1%	8.7%	3.8%		0.0%	3.8%
M6	6.8%	10.1%	0.9%		0.0%	0.9%
M7	7.7%	11.2%	2.5%		0.0%	2.5%
M8	8.2%	13.2%	3.1%		0.0%	3.1%
M9	8.6%	14.5%	4.6%		0.0%	4.6%
M10	9.5%	15.3%	6.7%		0.0%	6.7%
Total			**100.0%**	**100.0%**	**100.0%**	**0.0%**
Expected Return			6.6%	6.1%	5.8%	0.8%
Expected Risk			10.0%	10.5%	9.7%	2.0%
Efficiency Ratio						0.43

(continued)

EXHIBIT 5 *(continued)*

SAA with and without Hedge Funds—A Hypothetical Example

Panel D: Visual Representation of Panel C

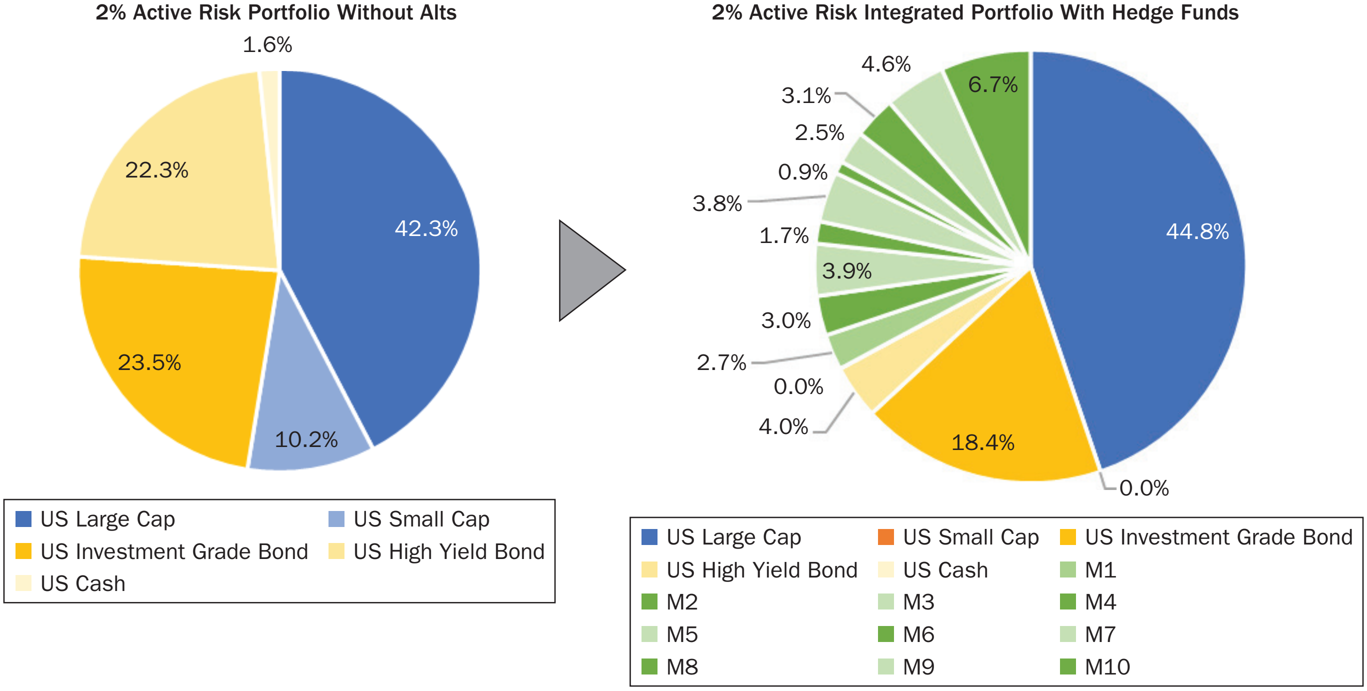

high-yield bonds. The crucial note to make here is this outcome is not set in stone. Instead, *it is very much a function of our capital market assumptions* and a delicate interplay between expected risk-adjusted returns coming from the core asset class exposures versus uncorrelated active returns. Core asset class exposures are contributed by the long-only part of the portfolio and hedge funds, combined. Uncorrelated active returns come from the hedge funds alone.

ALPHA/BETA SEPARATION IN THE PRESENCE OF LAGGED BETAS

Academic and practitioner communities alike broadly accept that many alternative asset classes exhibit *stale pricing* patterns that manifest themselves in a serial correlation of returns, or *lagged betas*. In the hedge fund space, such behavior is less prevalent but can be found in strategies exposed to less liquid asset-backed securities or special situations. In private equity and direct real estate spaces, this phenomenon is widespread, has been noted by many researchers (see, for example, Anson 2002 and Geltner 1991), and is linked to the market practice of obtaining valuations of private companies through self-reporting and/or appraisals. Methodologically, there is a broad body of literature on how to deal with stale pricing, which generally recommends to *unsmooth* the return time series before use. We would like to highlight three contributions that are particularly relevant.

As early as 2003, Conner (2003) noted that directly using private equities, illiquid hedge fund strategies, or other *smooth return* time series is misleading, as serial correlation hides the volatility of the underlying *true economic process* and obscures relationships with other liquid assets in the portfolio. Conner (2003) then suggested using a methodology from the real estate world (Geltner 1991) that estimates actual volatility by unsmoothing the time series of returns.

Pedersen, Page, and He (2014) went one step further and noted that in addition to needing to be unsmoothed, reported return time series ought to be analyzed through the lens of relevant factor exposures.

Finally, Rudin et al. (2019) introduced an econometric model that was philosophically similar to ones used by Pedersen, Page, and He (2014) and Conner (2003), but contained some enhancements improving the robustness of the parameter estimation process.

Let us denote the observed smooth return of an alternative asset as $r_t^{OBSERVED}$. The unsmoothing approach starts by recognizing the link between this return time series and the one for an unobserved true economic process return r_t^{TRUE}. Both Conner (2003) and Pedersen, Page, and He (2014) linked the two in the following way:

$$r_t^{OBSERVED} = \sum_{j=0}^{Q} w_j\, r_{t-j}^{TRUE} \tag{3}$$

Here Q is the maximum number of lags, and weights w_j sum up to one. Such formulation—together with empirical data—allowed both authors to estimate weights w_j and subsequently create a time series for true returns and evaluate its risk properties.

Additionally, Pedersen, Page, and He (2014) noted that the true economic process may be expressed as a linear combination of factor exposures denoted as F_t^K and idiosyncratic risks:

$$r_t^{TRUE} = \alpha^{TRUE} + \sum_{K=1}^{N} \beta^K F_t^K + \varepsilon_t^{TRUE} \tag{4}$$

Substituting Equation 4 into Equation 3 allows an estimation of factor exposures in the presence of smooth returns for a range of alternative strategies, completing the risk estimation process.

Although extremely helpful, this econometric model carried two drawbacks. First, Pedersen's process created a regression equation that contained an autocorrelated error term, requiring a Newey–West procedure to correct for it. This was rather inconvenient.

Additionally, Pedersen's method was a two-step procedure. First, they estimated the smoothness parameters w_j, and then estimated the factor betas β^K. Each estimation involved a regression and introduced a layer of estimation errors. If there were a way to estimate all parameters simultaneously, in a single regression, it could be helpful.

Rudin et al. (2019) pointed out that both drawbacks can be addressed if we start from an alternative but equivalent version of Equation 3 borrowed from the real estate literature (see Stefek and Suryanarayanan 2012):

$$r_t^{OBSERVED} = \theta_0\, r_t^{TRUE} + \sum_{j=1}^{Q} \theta_j\, r_{t-j}^{OBSERVED} \tag{5}$$

Here θ_j are unknown coefficients that can be estimated through regression analysis. Equation 5 is equivalent to Equation 3 in the sense that knowing θ_j, one can calculate w_j, and vice versa (also, $\sum_{j=0}^{Q} \theta_j = 1$). Equation 5 also has two material benefits. First, as long as one knows θ_j, true economic returns can be easily expressed through known observed returns by simply rearranging Equation 5:

$$r_t^{TRUE} = \frac{r_t^{OBSERVED} - \sum_{j=1}^{Q} \theta_j r_{t-j}^{OBSERVED}}{\theta_0} \tag{6}$$

Second, substituting Equation 4 into Equation 5 yields a regression equation with an error term that is *not* autocorrelated:

$$r_t^{OBSERVED} = \alpha + \left(1 - \sum_{j=1}^{Q} \theta_j\right) \sum_{K=1}^{N} \beta^K F_t^K + \sum_{j=1}^{Q} \theta_j r_{t-j}^{OBSERVED} + \varepsilon_t \tag{7}$$

Despite seeming complexity, Equation 7 allows—through a straightforward regression analysis—a simultaneous estimation of smoothing parameters θ_j and factor betas β^K.

Exhibit 6 illustrates the results of this approach applied to a commercial MBS bond index that happens to have a very pronounced serial correlation. Although not

EXHIBIT 6

Example of Alpha/Beta Separation in the Presence of Lagged Betas

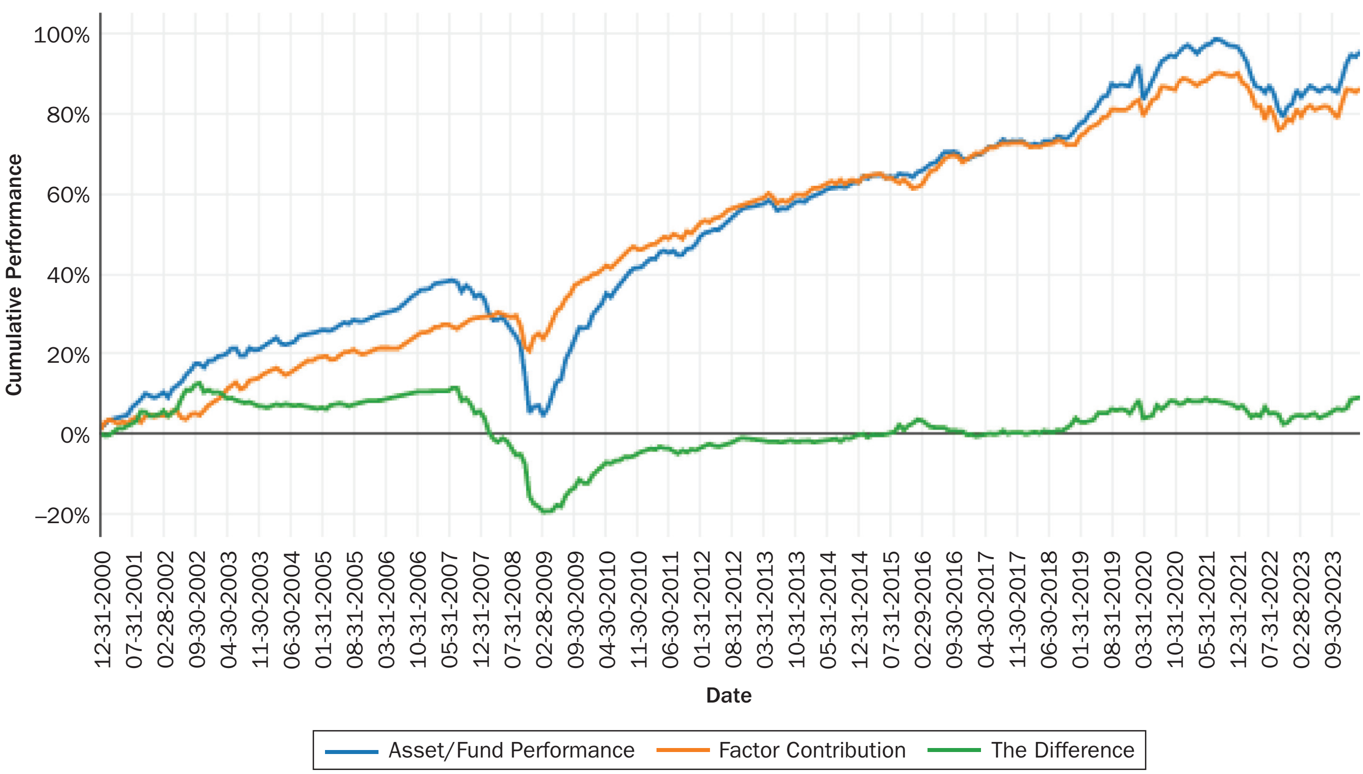

Factor Exposures		
Factor	**Beta**	**T Stat.**
US Treasuries	0.28	8.3
US High Yield	0.25	15.2
Lagged Asset Return		
Lag	**Theta**	**T Stat.**
1	0.27	6.3
2	0.08	1.9

a hedge fund strategy, this asset class is certainly an alternative when compared to more traditional Treasury, agency MBS, and corporate bonds. Exhibit 6 demonstrates that returns of this asset class can be well explained by exposures to US Treasuries and US high yield. It is also characterized by a material level of serial correlation (two lags were included; using the third lag didn't improve the replication power).

The general takeaway from this section is that SAA ought to reflect true risk properties of underlying assets. Lagged betas and factor exposures are key parts of the picture. As we will see in the next section, that remains true for assets such as private equity and private credit. That being said, those novel assets bring their own complexities practitioners need to deal with.

ALPHA VS. BETA IN PRIVATE EQUITY AND PRIVATE CREDIT

Private equity and private credit investments have become core components of many institutional portfolios. This is understandable—if viewed through the lens of risk-adjusted quarterly performance figures, private equity and credit have significant advantage over their public counterparts. This empirical fact (for example, historical information ratios of popular private equity indexes are roughly twice the one for the S&P 500) makes private equity and credit dominate publicly traded assets in most mean–variance-based total portfolio constructs. Institutional investors who are unwilling to follow suit have no choice but to either explicitly cap private allocations or artificially (and somewhat arbitrarily) penalize private equity expected return or risk figures.

Attempts to directly challenge quarterly performance figures as not linked directly with investors' experience over the lifetime of their investments so far have failed. For example, focused studies of the econometric relationship between the quarterly private and public mark-to-market returns (Conner 2003; Pedersen, Page, and He 2014; Rudin et al. 2019) seemed to confirm that even after the adjustments made for return smoothing, private assets enjoy significant alpha over their public counterparts.

At the same time, another substantial body of research (Phalippou 2014; L'Her et al. 2016; Ilmanen, Chandra, and McQuinn 2020) approached the relationship between private and public equity returns from a different, long-term (10-plus years) perspective. Using non-regression techniques (public market equivalents, discounted cash flow models, etc.), these studies found that over that long time horizon, the performance of private equity is generally in line with the performance of a leveraged public equity portfolio, with no material alpha in sight.

Contradictory outcomes of these two bodies of work were universally viewed as puzzling until Rudin and Farley (2022a) explained how to resolve this conundrum. Both bodies of work are correct, actually, but relate to very different investment horizons. Over the short term (less than a year), private asset returns are superior to public ones on a risk-adjusted basis, whereas over the long term (3-plus years), this advantage largely dissipates for an average private asset manager.

This newly articulated time horizon dimension of the asset allocation problem is somewhat novel but can be easily incorporated into the same integrated SAA alongside traditional long-only assets and hedge funds. Before turning to that, we will briefly summarize Rudin and Farley (2022a).

The authors started by challenging the traditional modeling approach schematically pictured as:

$$\begin{aligned}&\mathit{Observed\ private\ equity\ returns} \rightarrow \mathit{Unsmoothed\ private\ equity\ returns}\\&\mathit{Unsmoothed\ private\ equity\ returns} = \mathit{Alpha} + \mathit{Beta} \times \mathit{Public\ equity\ returns}\\&\qquad\qquad + \mathit{Residual}\end{aligned} \tag{8}$$

This approach implies that private equity is an alpha-generating complex strategy that is *connected* to public markets *but genuinely different*. But is it? Ever since Robert Shiller came up with the concept of *excess volatility* (Shiller 1981), it has been well accepted that public equity gyrations are driven partly by earnings expectations and partly by psychological phenomena (risk aversion, flows, momentum, etc.) that he collectively called excess volatility. We can schematically express this relationship as:

$$\textit{Public equity returns} = \textit{Change in economic fundamentals} + \textit{Excess volatility} \tag{9}$$

Private equity valuations processes, on the other hand, are primarily designed to mirror earnings expectations of underlying companies. Excess volatility that dominates public markets in the short term does not enter the private equity valuations in a material way, which may be the reason for why private equity betas and volatility are lower than comparable public benchmarks. Conversely, over the very long term, excess volatility washes away and economic fundamentals (earnings) are the only thing that truly matters. If one also reasonably assumes that underlying economic fundamentals of companies are the same, regardless of whether these companies are publicly or privately traded, it leads to a general expectation that long-term results of public and private equity investing ought to be similar to each other after a correction for leverage.

These qualitative considerations could be translated into the following schematic relationship between observed public and private equity prices:

$$\begin{aligned} &\textit{Observed private equity returns} \rightarrow \textit{Unsmoothed private equity returns} \\ &\textit{Public Equity Returns} = \textit{Alpha} + \frac{\textit{Unsmoothed Private Equity Returns}}{\textit{Leverage Advantage Ratio}} \\ &\qquad\qquad + \textit{Excess Volatility} \end{aligned} \tag{10}$$

where leverage advantage ratio (LAR) is simply a ratio of financial leverage of private versus public companies.

There is a material change in economic logic when we switch from Equation 8 to Equation 10. In the former, the traditional formulation public equity was considered thoroughly understood and hence used as an explanatory variable for an opaque asset such as private equity. In the new formulation of Rudin and Farley (2022a), private equity price changes are viewed as a relatively objective and relatively transparent window into market perceptions of economic fundamentals. In turn, public equities are a combination of these same fundamentals and something opaque and unpredictable—excess volatility.

Let us now proceed to a formal econometric model based on Equation 10. We will still follow Rudin and Farley (2022a), but omit details. As a first step, we constructed the unobserved true economic process return $r_t^{PE\_TRUE}$ time series using Equation 6. We chose a single lag model with ($\theta_1 = \theta$) as the most effective.

Now our informal model of the form of Equation 10 translates into the following:

$$r_t^{PUBLIC} = \alpha + \frac{r_t^{PE\_TRUE}}{LAR} + \varepsilon_t^{EXCESS} \tag{11}$$

where α represents the excess returns net of fees and the cost of leverage.

To make the problem tractable, authors assumed that both α and *LAR* are constant over time. This simplification is warranted given that we are interested in the long-term, multi-year relationships, and over these periods, the effects of both can be approximated by a constant.

One could assign a value to such a ratio based on information available to us (from a private equity manager, for example) or attempt to find this value by fitting empirical data. Rudin and Farley (2022a) used the fitting approach. We then calculated what we called the adjusted private equity returns:

$$r_t^{PE\_ADJUSTED} = \frac{1}{LAR} \frac{r_t^{PE\_OBSERVED} - \theta r_{t-1}^{PE\_OBSERVED}}{(1-\theta)} \tag{12}$$

These adjusted returns constitute what we believe is a fair comparison to public equity returns since all three elements—quarterly smoothing, leverage advantage, and presence of excess volatility—have been taken into account.

Exhibit 7 (from Rudin and Farley 2022a) summarizes the outcome of the approach. Both the smoothing coefficient θ and LAR were highly significant. On the other hand, alpha (α) that is supposed to measure differentiation between public and adjusted private equity returns was small and insignificant.

Exhibit 7 visualizes the main point Rudin and Farley (2022a) make: Using observed private equity returns as a proxy for the fundamentals component of public equity is highly effective and consistent with empirical evidence. This viewpoint rationalizes the gap between the observed private and public equity returns, resolving the long-standing conundrum. Over the long term, they evolve hand in hand, in line with observations made by Phalippou (2014), L'Her et al. (2016), and Ilmanen, Chandra, and McQuinn (2020). Over the short term, however, public equities are exposed to excess volatility. This makes public returns substantially inferior to private ones—on a risk-adjusted basis—over that time frame, as observed by Conner (2003), Pedersen, Page, and He (2014), Rudin et al. (2019), and many others.

Similar results were obtained for private credit when it was compared with the publicly traded high yield (Rudin and Farley 2022b). The leverage adjustment obtained

EXHIBIT 7

Public vs. Adjusted Private Equity Returns

	Estimate Value (T-statistic)
θ	0.385 (3.8)
LAR	1.323 (9.86)
α	–0.004 (–0.15)

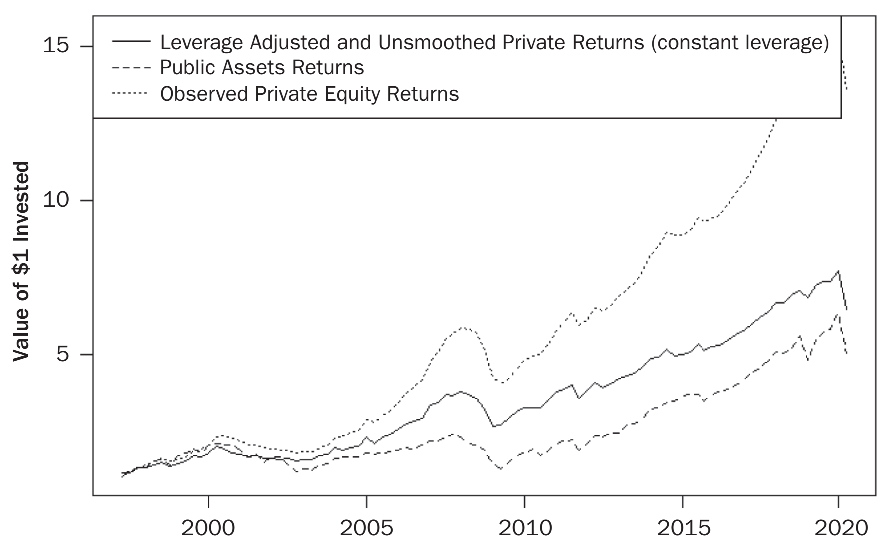

from fitting historical time series of private credit returns was estimated to be around $LAR = 1.6$.

We conclude this section with another practical benefit of the described approach. It allows us to directly link[5] long-term expected returns of private equity and private credit to those of their public counterparts:

$$\widehat{r^{PRIVATE}} = LAR\ \widehat{r^{PUBLIC}} + \alpha \tag{13}$$

Studies in Rudin and Farley (2022a, 2022b) suggest that α is small for an average manager, whereas *LAR* is 1.3 and 1.6 for an average private equity and private credit manager, respectively.

COMPARING RISKS OF PRIVATE AND PUBLIC ASSETS—AN INVESTMENT HORIZON DIMENSION

Now that we established a structural link between private and public assets behavior, we can turn to risk comparisons between the two.

Most institutional allocators use monthly or quarterly performance data to assess long-term risks of financial assets. A typical process involves taking monthly performance figures over a sufficiently long period of time (20 or 30 years), calculating monthly asset volatilities and correlations over these periods, annualizing volatilities by multiplying monthly figures by the square root of 12 (while leaving correlation estimates untouched), and then deploying results as *long-term risk estimates*.

This approach works well in the context of estimating asset risks over short investment horizons of, say, a month or a quarter; previously cited works by Conner (2003), Pedersen, Page, and He (2014), and Rudin et al. (2019) are examples. But in the context of strategic asset allocation, lengthening time horizons of our risk estimation would make a lot of sense. One way to go about this would be to decompose historical asset price patterns into persistent components and short-term noise and then focus on the former.

Exhibit 8 performs this decomposition on an example of S&P 500. We used a well-known procedure introduced by Hodrick and Prescott (1980) that separates total return time series into a combination of *persistent* and *transient* components. For S&P 500, those components differ not only in terms of speed of change but also in terms of their long-term dynamic. The slow, persistent component grows over time, reflecting growth in real economy and corporate earnings. The fast, transient component is directionless and strongly mean reverting. Granted, some of those characteristics were effectively ensured by our decomposition methodology, but they also reflect the dual nature of the price dynamic of most financial assets, including public equity and private credit.

In our quest to properly include broad assets into asset allocation, what would be the best way to evaluate risk characteristics of those assets? *The answer strongly*

[5] We caveat this broad statement with a word of caution. In contrast to the world of publicly traded assets in which one can easily invest passively into a broad swath of the market (for example, US equities), private investment exposure is always active and idiosyncratic. Practically investable products may substantially differ from uninvestable indexes and from each other (including differences in terms of leverage, strategic composition, alpha generation capability of managers [proven or expected], concentration, etc.). Articulating these differentiation points for specific products is beyond the scope of this article, yet these points may lead to substantial differentiation in both the return and risk components of such forecasts. Our recent research on the effects of concentration on both short- and long-term measures of private equity performance may serve as a case in point (Rudin et al. 2019).

EXHIBIT 8

S&P 500 Historical Price Time Series Split into Persistent and Transient Components Using the Hodrick–Prescott Filter

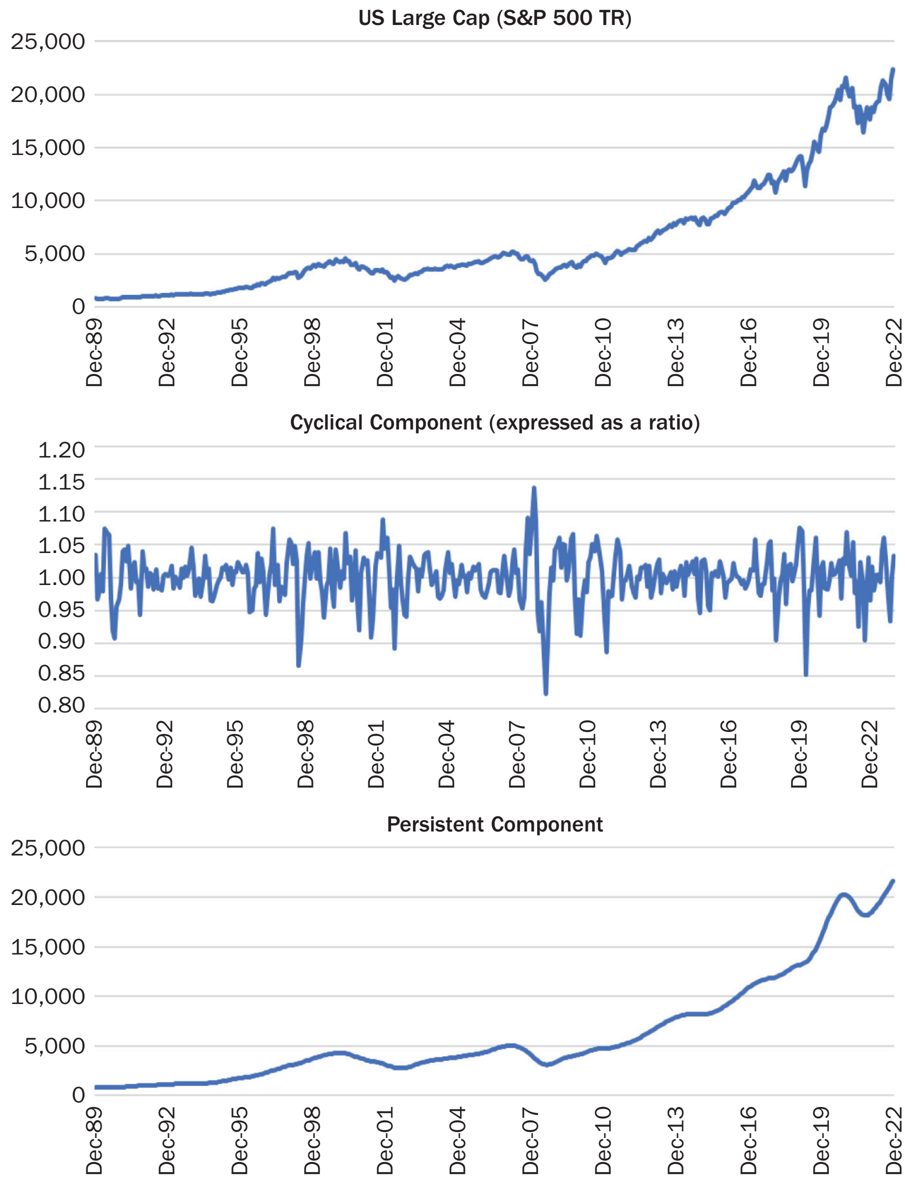

NOTES: The persistent component is created by minimizing both the mean squared deviation of the input signal from the persistent line and the curvature of it. The transient component is simply the difference between the historical price time series and the persistent component.

depends on our investment horizon. For short horizons, an asset's risk is best measured directly from its total price time series, which, in turn, is dominated by the transient component (i.e., noise).

As time horizon extends, however, using the persistent component as a starting point for risk estimation makes more sense. Noise dominates short-term price volatility, but over the longer horizons, this excess volatility subsides. For *strategic* asset allocation work, the investment horizon needs to be made appropriately long, and extracting the comparably persistent component from asset returns is well advised.

Although we chose public equities (S&P 500) as an introduction to this topic, the majority of, if not all, core asset classes experience similar risk dynamic (see Rudin and Farley (2022b) for a literature review). Over the long-term, asset prices are often anchored to some sort of slow-moving process,[6] but in the short-term, these same prices quasi-randomly cycle around their anchors.

Exhibit 9 shows annualized long-horizon volatility estimates for some of the major asset classes. Long-horizon volatility is lower than observed volatility for all assets, but the degree of reduction varies across asset classes. In a pure mean–variance portfolio optimization process, switching to long-horizon risk estimates would give a moderate boost to large-cap and global equities and US and UK sovereign bonds, at the expense of emerging market equities, non-US sovereign bonds, credit, and publicly traded alternatives.

Private assets also demonstrate a reduction in risk in the long horizon but to a much lesser degree. As a result, although observed near-term volatilities for private equity and private debt are lower than those for their public counterparts, the long-horizon volatilities are actually higher. *This is entirely intuitive.* As we discussed earlier, private equity and private debt are—over the long term—similar to leveraged versions of public equity and HY debt (at the rate of 1.3x and 1.6x, respectively). This is why when excess volatility dissipates, private equity and private debt emerge as riskier than with their public counterparts, despite their short-horizon risk figures telling a different story.

EXHIBIT 9

Long-Horizon Volatility Estimates for Major Asset Classes, January 2004–December 2023

Asset Class	Asset	Observed Price Vol.	Long Horizon Vol.	Vol. Ratio
Equity	US Large Cap	15.6%	4.6%	0.29
	Global Equities (ACWI) ex US	13.9%	4.7%	0.34
	Global Equities (ACWI)	14.5%	4.6%	0.31
	Europe	14.4%	4.6%	0.32
	Emerging Markets (EM)	17.1%	5.7%	0.34
Sovereign Bonds	US Government Bond	5.2%	1.4%	0.27
	Non-US Government Bonds	4.3%	1.3%	0.31
	UK Government Bonds	8.8%	2.4%	0.28
	Japanese Government Bonds	2.9%	0.8%	0.27
Credit	US Investment Grade Bond	5.0%	1.5%	0.30
	US High Yield Bond	9.6%	3.1%	0.33
Publicly Traded Alts	Hedge Funds	5.5%	2.2%	0.39
	Global Real Estate (REITs)	18.0%	6.5%	0.36
	Commodities	17.0%	5.7%	0.34
Private Alts	Private Equity	11.2%	7.5%	**0.67**
	Core Private Credit	7.7%	4.6%	**0.59**
	Opportunistic Private Credit	11.2%	6.9%	**0.62**
	Direct Real Estate—US	10.8%	9.0%	**0.83**

NOTES: We used the Hodrick–Prescott filter for this estimation. The observed volatility denotes the annualized standard deviation of monthly (for all categories except private alternatives) and quarterly (for private alternatives). Long-horizon volatility denotes the same calculation applied to the trend component of the Hodrick–Prescott filter output.

[6] One likely result of such dynamic would be for long-horizon asset volatility estimates to show better forecasting properties over multi-year periods than the short-horizon ones. And indeed, this behavior was recently confirmed empirically for a broad set of global equity and bond markets (Cardinale, Naik, and Sharma 2021).

EXHIBIT 10

Empirical Relationship between Private Equity Program Breadth and Its Idiosyncratic Risks

No. of Funds per Year	α	β	θ	Standard Deviation of the Residual
1	4.6%	0.53	0.39	15.6%
2				11.2%
3				9.4%
4				8.6%
5				8.3%
6				7.8%
7				7.4%
8				7.2%
9				7.0%
10				6.9%

Now, armed with the long-horizon risk estimation methodology, we are *almost* ready to proceed to our ultimate objective of building a self-consistent asset allocation framework. There is one additional complication to discuss—and that is the fact that popular private equity and credit indexes materially underestimate private programs' risks.

Most of the analysis of private equity returns done in academic literature use performance data for one of the broad private equity indexes. These contain thousands of individual fund contributions and are *uninvestable*. At the same time, a typical institutional private equity program only encompasses dozens (50–60 at most) of individual fund investments. Intuitively, it seems likely that a narrower program will carry higher idiosyncratic risks and thus higher volatility when compared to an index. An empirical study of synthetic private equity programs performed by Rudin et al. (2019) confirms and quantifies this intuition. By controlling the diversification level of private equity programs, the authors measured how idiosyncratic errors depend on those programs' breadth—see selected results in Exhibit 10. Idiosyncratic risks increase when a program's breadth narrows. Notably, this work was completed before the long-horizon concept was developed and uses unsmoothed quarterly returns for the study. That said, its general conclusion that program breadth matters should survive transition to the long-horizon risk estimation process.

We conclude this section with a summary recommendation on how to model private assets, such as private equity and credit, in broader portfolios. From the long-term expected return perspective, private assets could be modeled in consistent fashion with their public counterparts but amplified for the difference in leverage.

Risk picture is more complicated. Because private assets are much less sensitive to the short-term excess volatility than public assets, we recommend switching to long-horizon estimates for *all* assets, as this would bring everything to a level playing field. Also, comparatively higher concentration risks associated with typical private equity and credit programs suggest adding a (simulated or actual) idiosyncratic contribution to risk estimates for those assets. We will provide an example of how that can be done in the next section.

INTEGRATED SAA WITH PRIVATE ASSETS

In earlier sections, we established integrated SAA as a way to self-consistently add traditional and hedge fund assets to the same portfolio. Benefits of adding a long-horizon risk dimension to that framework are self-evident but seldom discussed in the literature. Delfim and Hoesli (2019) discussed the importance of long-horizon risk dimension for real estate. Outside of real estate applications, Rudin and Farley (2022b) built a multi-asset strategic portfolio construction framework that explicitly acknowledged term structure of risk expectations and, most importantly, suggested an actionable remedy. We now offer a very quick review of Rudin and Farley (2022b).

The authors pointed out that it is helpful to focus on the strategic optimality of SAA, even when only publicly traded assets are included. Benefits of this approach to SAA are even more pronounced in the presence of private assets. As the previous section demonstrated, in a long-horizon framework, risk-adjusted returns for private assets and related public counterparts are harmonized without the need to apply arbitrary adjustments to private asset risk figures.

Besides focusing on the long-horizon risk figures within the context of SAA, Rudin and Farley (2022b) also suggested incorporating a short-term tail risk dynamic in the form of additional constraints. The most straightforward way of doing this is in a form of historical *event risk constraints*. One can define past time periods that may be viewed as genuine tail events and then incorporate a series of constraints limiting portfolio losses over these periods (or average losses over multiple predefined periods). Tail risk constraint leads to a somewhat less efficient portfolio from the frontier perspective, but in exchange, such a portfolio would be better risk controlled over the short term.

The schematics of integrated SAA generally suitable for *all* types of (*typical*) alternatives is shown in Exhibit 11. It is similar to the original Exhibit 3 but incorporates all the innovations we have introduced so far: linking hedge fund and private assets to their long-only counterparts, explicit identification of factor exposures embedded in alternatives, Bayesian shrinkage applied to those components of alternative asset return expectations that are unreliable, and finally a dual-horizon approach to risk estimation that allows harmonization of risk across public and private assets.

Let us illustrate the process of integrated SAA by expanding our earlier case study (Exhibit 5) to include private equity. First, we need to move all risk calculations to the long-horizon space. To do that, we passed historical time series for all assets through a Hodrick–Prescott filter and used persistent components of the output for variance, covariance, and beta calculations. After we did that, we went through the same process described in an earlier case study (Exhibit 5). All alternative assets were viewed as a combination of something known (broad market exposures) and something unknown (their idiosyncratic contributions to return).

EXHIBIT 11

Schematics of Integrated SAA

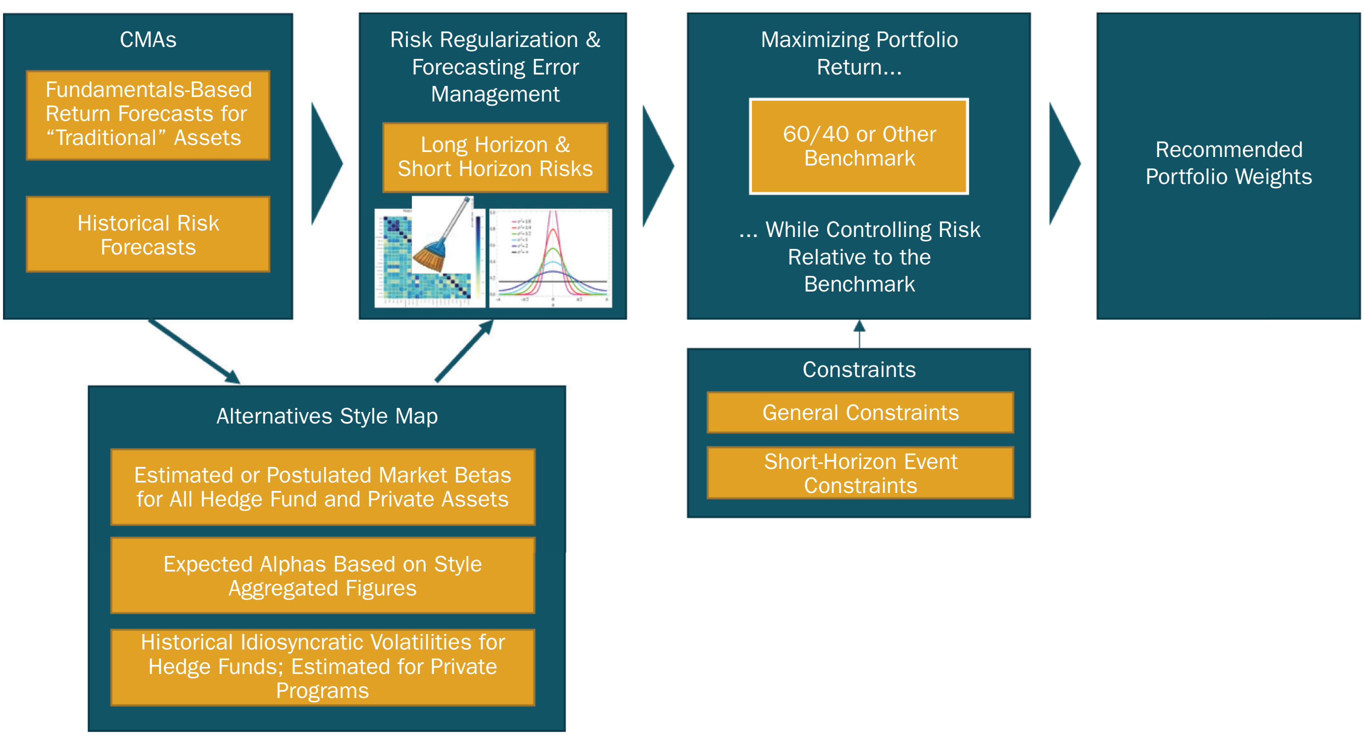

When it came to private equity, we modeled it returns as 1.3 × S&P 500 TR + Idiosyncratic risk. This is consistent with our calculated LAR of 1.3, alpha of zero associated with an average private equity fund and presence of idiosyncratic volatility. For the latter, we chose a level consistent with a midsize private equity program,[7] about 5.7% on the annualized basis.

Results are presented in Exhibit 12. We considered two cases. In the first (see Panels A and B), we deployed a pure long-horizon SAA methodology, and in the second (see Panels C and D), we added a short-term event risk constraint focused on the COVID-19 period of January 2020–March 2020, in accordance with a *dual-horizon* SAA philosophy.

Let us briefly walk through the results in Exhibit 12 and compare them with Exhibit 5. Outside of the newly incorporated allocation to private equity, recommended allocations did not change drastically. This is because risk figures for long-only assets scale downward in approximately the same proportion when one transitions to long-horizon risk estimation (see Exhibit 9), so their usefulness in SAA is maintained. Private equity received an allocation given its additive risk–return properties, but this allocation was moderate in size and fit neatly alongside equity-linked hedge funds and long-only equities. This private equity allocation also moderately increased when we added an event risk constraint. This makes intuitive sense: Private equity drawdown in Q1 2020 was only about –10%, which compares favorably to US large-cap equities (–19.6%), US small-cap equities (–32.6%), and even US high yield (–13%).

EXHIBIT 12

Integrated SAA With Hedge Funds and Private Equity—Hypothetical Example

Panel A: Long-Horizon Approach Without Short-Term Risk Constraints

Asset Classes	Expected Return	Expected Risk	SAA With HF	SAA Without HF	Benchmark SAA	Active Weight
US Large Cap	7.0%	4.6%	43.2%	42.8%	60.0%	–16.8%
US Small Cap	8.0%	5.4%	0.0%	9.9%	0.0%	0.0%
US Investment Grade Bond	4.0%	1.4%	19.8%	24.9%	40.0%	–20.2%
US High Yield Bond	6.0%	3.3%	0.5%	22.4%	0.0%	0.5%
US Cash	2.0%	0.5%	0.0%	0.0%	0.0%	0.0%
M1	3.0%	2.2%	0.0%		0.0%	0.0%
M2	7.8%	2.9%	10.8%		0.0%	10.8%
M3	6.6%	2.9%	0.2%		0.0%	0.2%
M4	6.8%	2.6%	0.0%		0.0%	0.0%
M5	8.5%	3.2%	1.2%		0.0%	1.2%
M6	7.5%	2.8%	4.4%		0.0%	4.4%
M7	9.2%	3.5%	2.0%		0.0%	2.0%
M8	7.9%	3.3%	3.1%		0.0%	3.1%
M9	9.6%	4.1%	0.8%		0.0%	0.8%
M10	9.7%	4.4%	7.3%		0.0%	7.3%
Private_Equity	15.3%	7.1%	6.6%		0.0%	6.6%
Total			**100.0%**	**100.0%**	**100.0%**	**0.0%**
Expected Return			7.4%	6.1%	5.8%	1.6%
Expected Risk			3.1%	3.2%	2.9%	0.6%
Efficiency Ratio			2.38	1.93	1.99	2.59

(continued)

[7] To get to that figure, we assumed that our private equity program allocates to five funds per year. Exhibit 10 suggests idiosyncratic risks of about 8.3% for such a program, which we consequently discounted to 0.69 × 8.3% = 5.7% using the long horizon to observe the volatility reduction ratio for private equity from Exhibit 9.

EXHIBIT 12 *(continued)*

Integrated SAA With Hedge Funds and Private Equity—Hypothetical Example

Panel B: Visual Representation of Panal A

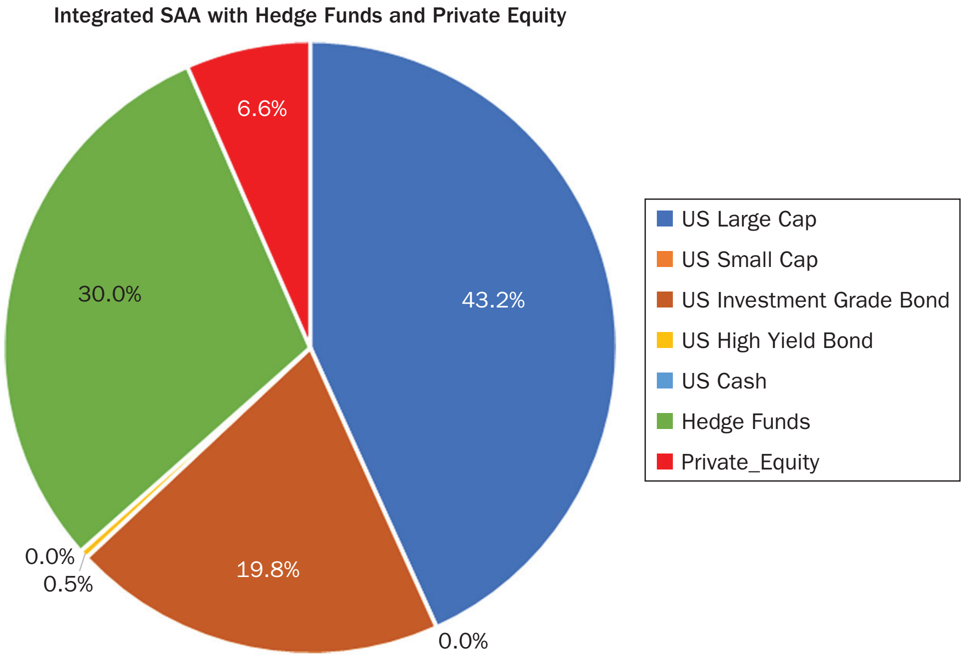

Panel C: Long-Horizon Approach with an Event Risk Constraint Limiting Hypothetical Portfolio Losses Observed During Q1 2020 to 300 bps (relative to the benchmark)

Asset Classes	Expected Return	Expected Risk	SAA With HF	SAA Without HF	Benchmark SAA	Active Weight
US Large Cap	7.0%	4.6%	38.8%	42.8%	60.0%	–21.2%
US Small Cap	8.0%	5.4%	0.0%	9.9%	0.0%	0.0%
US Investment Grade Bond	4.0%	1.4%	30.5%	24.9%	40.0%	–9.5%
US High Yield Bond	6.0%	3.3%	0.0%	22.4%	0.0%	0.0%
US Cash	2.0%	0.5%	0.0%	0.0%	0.0%	0.0%
M1	3.0%	2.2%	0.1%		0.0%	0.1%
M2	7.8%	2.9%	5.0%		0.0%	5.0%
M3	6.6%	2.9%	0.1%		0.0%	0.1%
M4	6.8%	2.6%	0.0%		0.0%	0.0%
M5	8.5%	3.2%	2.6%		0.0%	2.6%
M6	7.5%	2.8%	1.1%		0.0%	1.1%
M7	9.2%	3.5%	3.4%		0.0%	3.4%
M8	7.9%	3.3%	0.2%		0.0%	0.2%
M9	9.6%	4.1%	0.0%		0.0%	0.0%
M10	9.7%	4.4%	9.8%		0.0%	9.8%
Private_Equity	15.3%	7.1%	8.5%		0.0%	8.5%
Total			**100.0%**	**100.0%**	**100.0%**	**0.0%**
Expected Return			7.2%	6.1%	5.8%	1.4%
Expected Risk			2.9%	3.2%	2.9%	0.6%
Efficiency Ratio			2.49	1.93	1.99	2.43

(continued)

EXHIBIT 12 *(continued)*

Integrated SAA With Hedge Funds and Private Equity—Hypothetical Example

Panel D: Visual Representation of Panal C

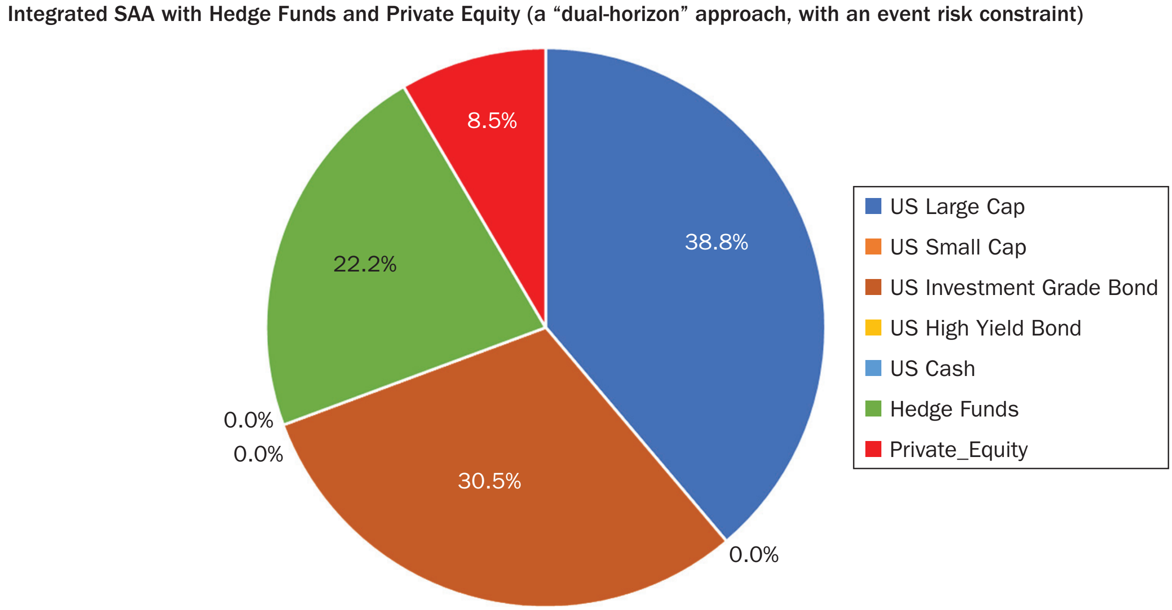

NOTE: The optimization targeted 60 bps in the long-horizon tracking error (this is roughly consistent with the 200 bps we used in Exhibit 5 when compared with the benchmark risk—10% annualized in Exhibit 5 and 3% annualized in the long-horizon calculation).

In the end, private equity received a solid but not overwhelming allocation of 8.5%, and that was without any artificial private equity allocation caps, special (upward) volatility adjustments, liquidity caps, or other extreme measures sometimes taken by asset allocators to keep private equity from dominating over other assets.

CONCLUDING REMARKS

Quantitative techniques for incorporating alternative investments and particularly private assets into strategic asset allocation remain unsettled and insufficiently covered in the literature. Hopefully, this article contributes to bridging this gap. We place all types of assets—from long-only to hedge funds to private equity and credit—on a level playing field by recognizing commonalities in their return drivers on one hand and treating idiosyncratic elements of those same returns in a differentiated way appropriate for each asset class.

We call our approach integrated SAA, and we see multiple advantages in deploying it for practical asset allocation tasks. It is self-consistent and broadly harmonizes asset allocation assumptions with empirical evidence. By separating persistent from transient in asset returns and explicitly focusing on the former, it adds clarity to the process. It also allows investors to understand the drivers of risk in their portfolio to properly assess potential opportunity costs required to manage both the long-term and the short-term risks appropriately.

We finish this article by pointing out that although an integrated approach to SAA was born out of our desire to build mixed portfolios that include alternative investments, it makes a lot of sense for purely liquid, long-only portfolios, as well. Using long-horizon volatility estimates for all publicly traded assets dampens transient effects of excess volatility and helps create a better strategic risk–return profile.

This method also promises to be more stable, as long-horizon volatility estimates have a lower forecasting error. Finally, balancing long-term optimality with short-term tail risk control is yet another rational objective for most institutional investors; our suggested framework offers a practical path toward achieving this objective through event constraints.

ACKNOWLEDGMENTS

We thank Pravesh Kumar, Nitin Singh, SK. Noor Mohammad, Vikas Mor, and Rajeeb Bharali for their help with putting together the case studies and other exhibits in this article, as well as for their valuable discussions.

REFERENCES

Anson, M. 2002. "Managed Pricing and the Rule of Conservatism in Private Equity Portfolios." *The Journal of Private Equity* 5 (2): 18–30.

Brandt, M. W. 2010. "Portfolio Choice Problems." In *Handbook of Financial Econometrics, Volume 1: Tools and Techniques*, edited by Y. At-Sahalia and L.P. Hansen, pp. 270–330. Amsterdam: North Holland.

Bun, J., J. P. Bouchaud, and M. Potters. 2016. "Cleaning Correlation Matrices." Risk.net.

Cardinale, M., N. Y. Naik, and V. Sharma. 2021. "Forecasting Long-Horizon Volatility for Strategic Asset Allocation." *The Journal of Portfolio Management* 47 (4): 83–98.

Conner, A. 2003. "Asset Allocation Effects of Adjusting Alternative Assets for Stale Pricing." *The Journal of Alternative Investments* 6 (3): 42–52.

Delfim, J.-C., and M. Hoesli. 2019. "Real Estate in Mixed-Asset Portfolios for Various Investment Horizons." *The Journal of Portfolio Management* 45 (7): 141–158.

Geltner, D. 1991. "Smoothing in Appraisal-Based Returns." *Journal of Real Estate Finance and Economics* 4 (3): 327–345.

Hodrick, R. J., and E. C. Prescott. 1980. "Postwar US Business Cycles: An Empirical Investigation." Working paper, Carnegie Mellon University.

Ilmanen, A., S. Chandra, and N. McQuinn. 2020. "Demystifying Illiquid Assets: Expected Returns for Private Equity." *The Journal of Alternative Investments* 22 (3): 8–22.

Lhabitant, F.-S. 2004. *Hedge Funds: Quantitative Insights*. Hoboken: John Wiley & Sons.

L'Her, J.-F., R. Stoyanova, K. Shaw, W. Scott, and C. Lai. 2016. "A Bottom-Up Approach to the Risk-Adjusted Performance of the Buyout Fund Market." *Financial Analyst Journal* 72 (4): 36–48.

Molyboga, M., and L. E. Swedroe. 2023. *Your Essential Guide to Quantitative Hedge Fund Investing*. Boca Raton: Chapman and Hall.

Pedersen, N., S. Page, and F. He. 2014. "Asset Allocation: Risk Models for Alternative Investments." *Financial Analysts Journal* 70 (3): 34–45.

Phalippou, L. 2014. "Performance of Buyout Funds Revisited." *Review of Finance* 18 (1): 189–218.

Rudin, A. 2018. "Persistence of Hedge Fund Returns and Fee-Aware Portfolio Construction." *The Journal of Portfolio Management* 44 (5): 103–112.

Rudin, A., and D. Farley. 2022a. "Public and Private Equity Returns: Different or Same?" *The Journal of Portfolio Management* 48 (3): 117–127.

——. 2022b. "Dual-Horizon Strategic Asset Allocation." *The Journal of Portfolio Management* 48 (4): 59–72.

Rudin, A., J. Mao, N. R. Zhang, and A.-M. Fink. 2019. "Fitting Private Equity into the Total Portfolio Framework." *The Journal of Portfolio Management* 46 (1): 60–73.

Rudin, A., V. Mor, and D. Farley. 2020. "Adaptive Optimal Risk Budgeting." *The Journal of Portfolio Management* 46 (6): 147–158.

Shiller, R. J. 1981. "Do Stock Prices Move too Much to Be Justified by Subsequent Changes in Dividends?" *American Economic Review* 71 (3): 421–435.

Stefek, D., and R. Suryanarayanan. 2012. "Private and Public Real Estate: What Is the Link?" *The Journal of Alternative Investments* 14 (3): 66–75.

DISCLOSURES

Marketing Communication

ssga.com

Australia: State Street Global Advisors, Australia, Limited (ABN 42 003 914 225) is the holder of an Australian Financial Services Licence (AFSL Number 238276). Registered office: Level 14, 420 George Street, Sydney, NSW 2000, Australia T: +612 9240-7600. F: +612 9240-7611.

Belgium: State Street Global Advisors Belgium, Chaussée de La Hulpe 185, 1170 Brussels, Belgium. T: +32 2 663 2036. State Street Global Advisors Belgium is a branch office of State Street Global Advisors Europe Limited, registered in Ireland with company number 49934, authorised and regulated by the Central Bank of Ireland, and whose registered office is at 78 Sir John Rogerson's Quay, Dublin 2.

Canada: State Street Global Advisors, Ltd., 1981 McGill College Avenue, Suite 500 , Montreal, Quebec, H3A 3A8, T: +514 282 2400 and 30 Adelaide Street East Suite 1100, Toronto, Ontario M5C 3G6. T: +1647 775 5900.

Dubai: State Street Global Advisors Limited, DIFC branch is regulated by the Dubai Financial Services Authority (DFSA). This document is intended for Professional Clients or Market Counterparties only as defined by the DFSA and no other person should act upon it. State Street Global Advisors Limited, DIFC Branch, OT 01-39, 1st Floor, Central Park Towers, DIFC, P.O Box 507448, Dubai, United Arab Emirates. Regulated by the DFSA. Telephone: +971 4 871 9100.

France: State Street Global Advisors Europe Limited, France Branch ("State Street Global Advisors France") is a branch of State Street Global Advisors Europe Limited, registered in Ireland with company number 49934, authorised and regulated by the Central Bank of Ireland, and whose registered office is at 78 Sir John Rogerson's Quay, Dublin 2. State Street Global Advisors France is registered in France with company number RCS Nanterre 899 183 289, and its office is located at Coeur Défense - Tour A - La Défense 4, 33e étage, 100, Esplanade du Général de Gaulle, 92 931 Paris La Défense Cedex, France. T: +33 1 44 45 40 00. F: +33 1 44 45 41 92.

Germany: State Street Global Advisors Europe Limited, Branch in Germany, Brienner Strasse 59, D-80333 Munich, Germany with a representation office at Brussels Street 1-3, D-60327 Frankfurt am Main Germany ("State Street Global Advisors Germany"). Munich T +49 (0)89 55878 400. Frankfurt T +49 (0)69 667745 000. State Street Global Advisors Germany is a branch of State Street Global Advisors Europe Limited, registered in Ireland with company number 49934, authorised and regulated by the Central Bank of Ireland, and whose registered office is at 78 Sir John Rogerson's Quay, Dublin 2.

Hong Kong: State Street Global Advisors Asia Limited, 68/F, Two International Finance Centre, 8 Finance Street, Central, Hong Kong. T: +852 2103-0288. F: +852 2103-0200.

Ireland: State Street Global Advisors Europe Limited ("SSGAEL"), regulated by the Central Bank of Ireland. Registered office address 78 Sir John Rogerson's Quay, Dublin 2. Registered number 49934. T: +353 (0)1 776 3000. Fax: +353 (0)1 776 3300. Web: www.ssga.com.

Italy: State Street Global Advisors Europe Limited, Italy Branch ("State Street Global Advisors Italy") is a branch of State Street Global Advisors Europe Limited, registered in Ireland with company number 49934, authorised and regulated by the Central Bank of Ireland, and whose registered office is at 78 Sir John Rogerson's Quay, Dublin 2. State Street Global Advisors Italy is registered in Italy with company number 11871450968 - REA: 2628603 and VAT number 11871450968, and its office is located at Via Ferrante Aporti, 10 - 20125 Milan, Italy. T: +39 02 32066 1

8178281.1.1.GBL.INST

Exp. Date: 31/07/2030