from **asset allocation** *to* **risk allocation**

The Risk Allocation Framework

C|A CAMBRIDGE ASSOCIATES LLC

from **asset allocation** *to* **risk allocation**

The Risk Allocation Framework

List of Figures

Appendices

Executive Summary

- The most recent step in the evolution of portfolio construction practices has been a shift from an asset allocation–centered process to a more comprehensive risk allocation–based process. Cambridge Associates' Risk Allocation Framework considers multiple dimensions of risk and return trade-offs when building portfolios and evaluates the consequences of risk allocation decisions during normal and stressed markets.

- Investors have traditionally constructed portfolios by considering how much to allocate to different asset classes. Since the 1970s-style balanced fund allocated to domestic stocks and bonds, asset allocation has generally evolved as follows:

  1. Replacement of balanced portfolios by separate asset classes, including asset classes such as foreign market equities.

  2. Birth of the "style box" as investors hired separate managers to focus on large-cap, small-cap, growth, and value stocks.

  3. Broader adoption of real estate and other forms of private investments (e.g., venture capital and buyout funds).

  4. Addition of distressed securities, commodities, natural resources, and various kinds of hedge funds.

- Yet the problem became that several of these more recently introduced "asset classes" actually have common risk factors that cross "asset class" boundaries. Examples include equity risk in distressed securities and natural resources equities, and illiquidity risk in hedge funds and commingled funds—particularly in stressed environments. Thus it became increasingly difficult to recognize, without significant analysis, just how much equity risk (for example) might be embedded in a portfolio that owned lots of assets not named "equities."

- To clarify matters, investors increasingly have constructed portfolios on the basis of the role they expected different kinds of investments to play in the portfolio (i.e., *role-in-portfolio exposures*), even if they still allocated investments to traditional asset classes.

- The Risk Allocation Framework takes this evolution a step further by considering not only the role that different investments might play in the portfolio, but how and in what ways such investments contribute to or mitigate various forms of portfolio risk. The framework combines careful attention to risk allocation in the context of the risk sensitivities and limitations of a long-term investment portfolio (LTIP) given its role in the broader organization. Since risk exposures move over time, we monitor risk allocation and performance attribution dynamically.

- There are four main components to the Risk Allocation Framework.

  - *Enterprise Review:* Simplify and clarify the process for identifying the degree of dependence on and integration of the LTIP in the total enterprise and identifying any constraints on portfolio construction. This facilitates informed decision making about appropriate fundamental exposures to include in the portfolio.

  - *Policy Setting:* Set top-down objectives to clearly express investors' most timeless and fundamental risk tolerances and objectives. Policy reflects desired role-in-portfolio exposures (e.g., diversified growth, deflation hedge, and inflation sensitive), value-added performance objectives, and common risk factors of equity beta, illiquidity tolerance, and foreign currency risk. The policy is designed to provide those implementing portfolios with the appropriate guidelines for meeting long-term objectives.

  - *Implementation:* Determine allocations to most effectively implement investment policy objectives in the current environment. Implementation

includes all decisions that result in differences between the actual portfolio and the policy portfolio, including more detailed asset allocation, manager structure, and manager selection. Measure and monitor implementation decisions to make sure risks taken are consistent with conviction about their potential value added relative to policy. Given the changing nature of implementation risks and potential rewards, the framework uses a dynamic approach to understand risk characteristics of portfolio positions in the current environment and put them in a historical context.

- *Ongoing Performance Monitoring:* Use performance measurement and attribution analysis to understand the ways in which value has been added to (or subtracted from) the portfolio. Performance measurement serves as a feedback loop into continuous improvement in portfolio management.

- Skeptics might question whether the Risk Allocation Framework is an evolutionary step forward or merely just the same old fellow, dressed in a new suit. Asset allocation—informed by rigorous valuation analysis—and manager selection remain important parts of the Risk Allocation Framework, so there is some family resemblance. In fact, this is by design, as we preserved the best practices of traditional portfolio construction (often referred to as "the endowment model" because endowments were the earliest adopters).

- However, the Risk Allocation Framework's primary merits are that it enables investors first to more rigorously *construct* portfolios designed to realize their investment objectives and second to understand much more clearly how best to *manage* such a portfolio dynamically to improve the likelihood that it performs as anticipated. ■

The most recent step in the evolution of portfolio construction practices has been a shift from an asset allocation–centered process to a more comprehensive risk allocation–based process. Cambridge Associates' Risk Allocation Framework considers multiple dimensions of risk and return trade-offs when building portfolios and evaluates the consequences of risk allocation decisions during normal and stressed markets.

In this paper, we share our views on why asset allocation has developed into risk allocation and provide an introduction to the Risk Allocation Framework. We focus on the four steps of the portfolio construction process that have evolved the most to permit a more rigorous evaluation of varied risk and return trade-offs: the Enterprise Review, Policy Setting, Implementation, and Ongoing Performance Monitoring.

We describe the Risk Allocation Framework by using the example of an endowment in order to simplify the discussion. The framework can be modified to apply to virtually any investor, from families to pension fund sponsors to hospitals. (See Appendix A for a discussion on the application of the Risk Allocation Framework to other investors.)

The Evolutionary Process

Investors traditionally have constructed portfolios by considering how much to allocate to different asset classes. Since the 1970s-style balanced fund allocated to domestic stocks and bonds, asset allocation has generally evolved as follows:

1. Replacement of balanced portfolios by separate asset classes, including asset classes such as foreign market equities.
2. Birth of the "style box" as investors hired separate managers to focus on large-cap, small-cap, growth, and value stocks.
3. Broader adoption of real estate and other forms of private investments (e.g., venture capital and buyout funds).
4. Addition of distressed securities, commodities, natural resources, and various kinds of hedge funds.

Yet the problem became that several of these more recently introduced "asset classes" actually have common risk factors that cross "asset class" boundaries. Examples include equity risk in distressed securities and natural resources equities, and illiquidity risk in hedge funds and commingled funds—particularly in stressed environments. Thus it became increasingly difficult to recognize, without significant analysis, just how much equity risk (for example) might be embedded in a portfolio that owned lots of assets not named "equities." (See Figure 1.)

The proliferation of "asset classes" with common underlying risk factors can provide a false sense of security that portfolios are diversified.

Further, creating a diversified portfolio involves far more than just investing in many asset classes. The proliferation of "asset classes" with common underlying risk factors can provide a false sense of security that portfolios are diversified. For example, high-yield bonds and non-Agency mortgages are fixed income, but unlike sovereign bonds they cannot be expected to perform well during periods of economic contraction when credit spreads widen. (See Figure 2.)

To clarify matters, investors increasingly have constructed portfolios on the basis of the role they expected different kinds of investments to play in the portfolio (i.e., *role-in-portfolio exposures*), even if they still allocated investments to traditional asset classes.

Figure 1 **Different Portfolios with the Same Asset Class Allocation**

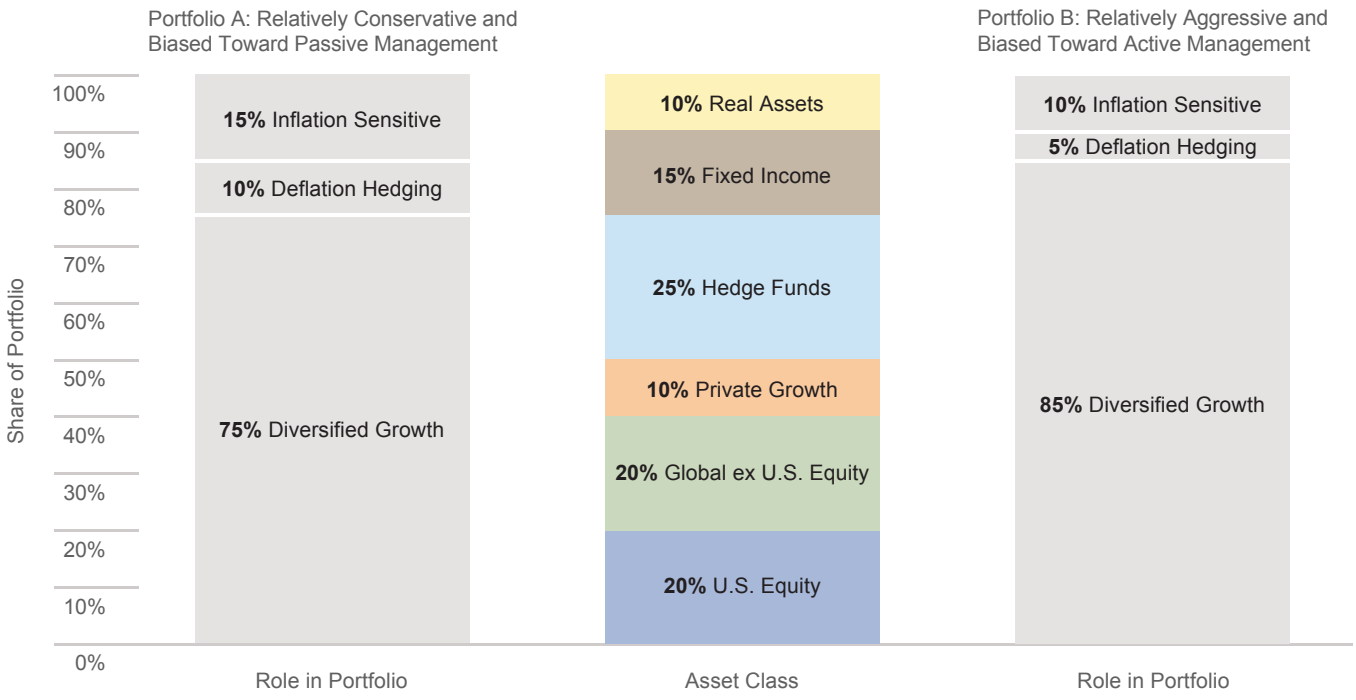

This exhibit shows how traditional asset class terminology can mask important differences in portfolios, both at the highest "role-in-portfolio" exposure level and at more detailed levels. These differences have implications for the expected return and volatility of both *policy* and *implementation*.

Above we show two U.S. investor portfolios, A and B, at a high level based on the roles investments serve in the portfolio. Their proportions of "diversified growth" (our classification for equity and equity-like investments), deflation hedge, and inflation-sensitive components differ. Portfolio A is allocated 75% to diversified growth, 10% to deflation hedge, and 15% to inflation sensitive, whereas Portfolio B takes 5% each from the macro hedges and adds that to diversified growth, resulting in an 85%-5%-10% allocation. This results in *policy* expectations of higher equity beta, higher volatility, and slightly higher compound return for Portfolio B than for Portfolio A.

The middle portfolio shows asset class allocations for both A and B, with the composites at a typical level of detail for asset allocation–based policy portfolios. Note that despite the different role-in-portfolio allocations, the asset class allocations are identical. How is this possible?

On the following page we show the same portfolios with detail down to manager structure. Note that the asset class term "fixed income" comprises assets that deliver rather diverse exposures and thus fall under different role-in-portfolio groups. Specifically, investment-grade and high-yield corporate bonds are more growth oriented and are part of diversified growth, inflation-linked bonds are part of inflation sensitive, and only high-quality sovereign bonds constitute deflation hedges. Portfolio A's 15% allocation to fixed income comprises 10% in Treasuries and 5% in TIPS, which may be viewed as a relatively conservative approach to fixed income. Yet Portfolio B's 15% allocation comprises 5% each in investment-grade and high-yield corporates, 5% in Treasuries, and none in TIPS, a reasonably common approach to seek higher returns at the expense of assets more protective against deflation.

(continued on the next page)

Figure 1 Different Portfolios with the Same Asset Class Allocation

(continued from the previous page)

Other asset class terms also comprise underlying investments with rather different exposure, return, and volatility profiles. The manager structure detail reveals further implementation differences between the portfolios that are not apparent at the asset class level. Contrasted with the relatively conservative Portfolio A, the more aggressive Portfolio B:

- in U.S. equities, includes more small-cap;
- in global ex U.S. equities, includes more emerging markets;
- in private growth, includes more venture capital;
- in hedge funds, tilts toward more long-biased long/short managers; and
- in real assets, favors private over public inflation-sensitive assets, increasing illiquidity.

Implementation could further differ at the manager selection level. In this example Portfolio A invests passively wherever possible, whereas Portfolio B invests actively and targets higher-beta vehicles, adding active risk but also the opportunity for value added.

In sum, Portfolio A's policy has a compound real expected return of 5% with 9.9% volatility, and as implemented, an expected return of 5.7%. Portfolio B's policy has an expected return of 5.1% with 11.1% volatility, and as implemented as much as a 6.7% return (based on reasonable assumptions for value added from manager selection active risk). All from two portfolios with the same asset class allocations!

The Risk Allocation Framework takes into account the limitations of asset allocation definitions illustrated in this example. The framework incorporates careful evaluation of risk sensitivities to integrate objectives and constraints along with role-in-portfolio allocation targets into investment policy. Policy setting includes factors such as expectations for portfolio equity beta, illiquidity tolerance, desired foreign currency risk, and value-added objectives in addition to return and volatility goals, all of which can differ meaningfully across portfolios with like asset allocation. And the framework couples this robust policy setting with rigorous dynamic evaluation of risk allocation and performance attribution.

Manager Structure

	Portfolio A	Portfolio B
Real Assets Private RE	2.5%	5.0%
Real Assets REITs	2.5%	-
Real Assets Private O&G	2.0%	3.0%
Real Assets Public NRE	2.0%	2.0%
Real Assets Commodities	1.0%	-
Fixed Income TIPS/ILBs	5.0%	-
Fixed Income Core Sovereign	10.0%	5.0%
Fixed Income Inv-Grade Corp	-	5.0%
Fixed Income High-Yield Bonds	-	5.0%
Hedge Funds Absolute Return	12.5%	7.5%
Hedge Funds Long/Short	12.5%	17.5%
Private Growth PE	6.5%	3.5%
Private Growth VC	3.5%	6.5%
Emerging Markets Equity	5.0%	10.0%
Developed ex U.S. Equity	15.0%	10.0%
U.S. Equity SCV	-	2.0%
U.S. Equity SCG	-	2.0%
U.S. Equity LCV	5.0%	5.0%
U.S. Equity LCG	5.0%	5.0%
U.S. Equity Core	10.0%	6.0%

Metrics

	Portfolio A	Portfolio B
Policy		
Real Expected AACR	5.0%	5.1%
Standard Deviation	9.9%	11.1%
Equity Beta (Portfolio)	0.65	0.69
Implementation		
Real AACR (ex Manager Selection)	5.7%	6.2%
Real AACR (incl Manager Selection)	n/a	6.7%
Value Add	70 bps	160 bps
Illiquidity	14.5%	18.0%

Figure 2 **Is Your Portfolio Diversified? Asset Allocation Is Only Part of the Story** *As of December 31, 2011*

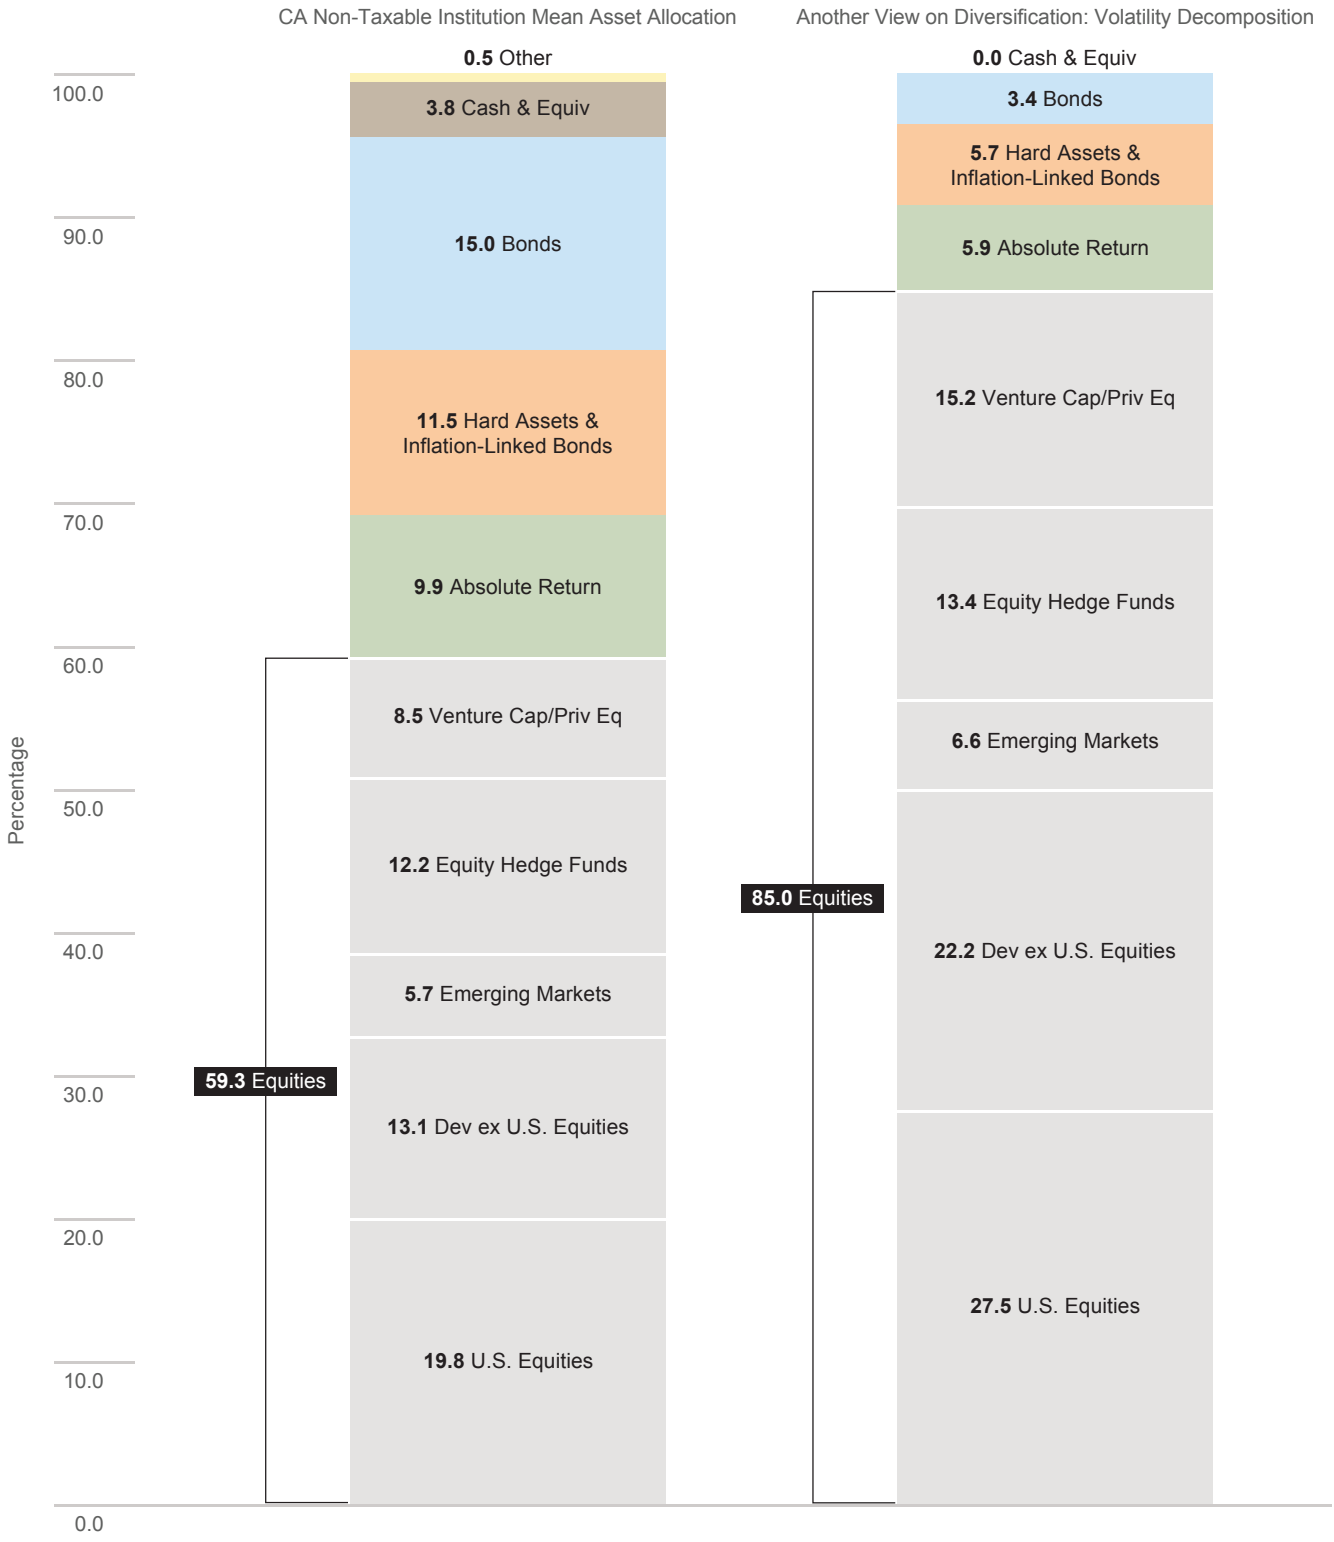

Notes: Universe represents all CA non-taxable institutions. For modeling purposes, "Other" is allocated 50/50 to "Equity Hedge Funds" and "Absolute Return." Percentages may not sum to 100% due to rounding.

Risk Allocation

The Risk Allocation Framework takes this evolution a step further by considering not only the role that different investments might play in the portfolio, but how and in what ways such investments contribute to or mitigate various forms of portfolio risk. The framework combines careful attention to risk allocation in the context of the risk sensitivities and limitations of a long-term investment portfolio (LTIP) given its role in the broader organization. Since risk exposures move over time, we monitor risk allocation and performance attribution dynamically.

The framework evaluates risk/return trade-offs to meet return objectives given capital markets history, current

Figure 3 **Considering the Trade-Offs in Risk Allocation**

GOAL	ACTIONS	BENEFIT	COST
Increase Return	Increase growth and/or "return enhancers" Reduce "diversifiers," low-beta growth engine, and/or macro-economic hedges	Purchasing power preservation Portfolio growth Spending growth Reputation/peer comparisons	Higher short-term portfolio volatility Higher spending volatility (dependent on spending rule) Potential for increased tail risk Potential for increased equity beta or illiquidity risk
Reduce Volatility	Increase "diversifiers," low-beta growth, and/or macroeconomic hedges Reduce growth and/or "return enhancers"	Lower risk of short-term decline Lower portfolio volatility Lower spending volatility	Reduced probability of purchasing power preservation Reduced portfolio growth Reduced spending growth Potential for increased complexity if investing in more "diversifiers"/hedge funds
Protect Against Macro Risk	Increase asset-backed macroeconomic hedges (e.g., high-quality sovereign bonds, hard assets) Institute portfolio/derivatives-based hedges	Lower portfolio declines during periods of stress Increased liquidity/dry powder during periods of stress for operations and/or for portfolio	Reduced probability of purchasing power preservation Reduced portfolio growth Reduced spending growth Behavioral risk associated with any hedging, especially derivatives-based hedging Basis risk associated with all hedging
Reduce Foreign Currency–Related Volatility	Hedge FX exposure	Potentially lower portfolio volatility	Operational complexity Cost of hedging Opportunity cost of holding increased liquidity for margin Behavioral risk
Increase Potential for Value Added	Employ active managers Engage in tactical asset allocation	Increase return Diversify sources of return	Increased tracking error (risk) versus policy Potential for lower returns due to uncertain value added, and costs of oversight Reduced transparency and liquidity from active managers Higher fees for active managers

conditions, and individual investor skills, resources, and risk tolerances (see Figure 3). For example:

- If you incorporate more pure[1] forms of diversification (e.g., high-quality, intermediate- to long-duration bonds, cash, commodities) you may get more protection during difficult environments, but you also reduce expected returns.
- If you take more active risk (defined as differences between the actual portfolio and the policy portfolio[2]) you increase expected returns if you have skill (or good luck!), but also increase potential downside risk.
- And while you pursue manager value added to improve returns, you may also take more illiquidity risk and more equity (or equity-like) risk. Value-added opportunities tend to be more plentiful in private investments and hedge funds, both of which also come with varying degrees of such risks.

We describe each of the four stages of the Risk Allocation Framework that has evolved to meet portfolio construction and management needs in the modern asset allocation environment. We also share our thoughts on governance, which reflect our long-standing views, but are important enough to merit inclusion in any discussion of best practices. (See Figure 4 at right for a high-level schematic of our portfolio construction process.)

Figure 4 **Risk Allocation Framework Roadmap**

Areas of focus for enhancements to our process are highlighted in blue.

1 GOVERNANCE	An overview of best practices in governance.
2 INVESTMENT PHILOSOPHY	An explanation of Cambridge Associates' investment approach, providing the background for subsequent sections.
3 ENTERPRISE REVIEW	An examination of enterprise-level budgetary, operating, and risk considerations to determine whether such considerations may constrain the construction of the long-term investment portfolio, and vice versa.
4 POLICY SETTING	Starting from the institution's specific goals and risk tolerance, a high-level, top-down set of objectives and constraints to set long-term expectations and guide portfolio implementation.
5 IMPLEMENTATION	Evaluation of active decisions embedded in the portfolio, from the recommended policy portfolio to the actual portfolio, into finer components. Review positions taken and evaluate their relative sizing and attractiveness in the current environment. Provides an *ex ante* assessment of risk.
6 MONITORING	Measurement of the value added from the different components of active management relative to carefully selected investible benchmarks. Guidelines for rebalancing.

[1] That is, assets without substantial exposures to factors you intend to diversify away from, such as equity beta.

[2] Specifically defined as the standard deviation of value-added returns, where value added is the difference between actual portfolio returns and benchmark returns. Active risk can be calculated for the total portfolio or any individual active position (e.g., manager performance in relation to the manager benchmark).

Consider the Context: The Enterprise Review

By starting with a process to identify any constraints on portfolio construction, the Risk Allocation Framework aims to correct an oversight that may occur in more limited examinations of investment policy: what role does the LTIP play in the institution's business model? You should not necessarily begin with the assumption that the LTIP is unconstrained in how it is invested—although that may indeed be the case for an institution where there is little functional linkage between the LTIP and the institution's budget and operations. More often, because there is a linkage, you must check whether there are non-investment-related circumstances that should be considered when developing the portfolio's risk profile. While this is important even in ordinary markets, the global financial crisis of 2008–09 revealed for many institutions the demands placed by their business models on their LTIP in stressed environments (and the reverse).

The ultimate purpose of the Enterprise Review is to uncover any issues that might be relevant, and thus reduce the chance of meaningful surprises, especially during a crisis.

The Enterprise Review is a comprehensive examination of financial circumstances, risk attitudes, and governance issues. The ultimate purpose is to uncover as many of these issues that might be relevant, and thus reduce the chance of meaningful surprises, especially during a crisis.

The key first area to be addressed is the bridge between the endowment portfolio and the institution's operations. From the LTIP's viewpoint, what is the spending (payout) policy? And from the perspective of the institution, what is the level of operating revenue dependence on such distributions? The former has implications for what the LTIP's return target should be, while the latter is more critical when thinking about the LTIP's volatility target.

Other issues to consider in this review include acquiring an understanding of:

- *operating liquidity*, in case an institutional need for liquidity during a stress period prompts endowment withdrawals above spending policy;
- *balance sheet inflexibilities*, e.g., whether unrestricted assets are a relatively low share of the total;
- *liabilities*, including amount of debt, pension, and capital commitments;
- *debt level and structure*, focusing particularly on any putable variable rate financing;
- *access to external liquidity*, including taxable debt and lines of credit;
- *cost structure*, e.g., whether there are high fixed costs and concomitant budgetary demands; and
- *revenue structure*, including whether economic environments that negatively impact the LTIP might have a similar effect on non-endowment revenue streams.

All of these considerations are designed to help identify implications for the investment of the LTIP. How much institutional liquidity is there, relative to both the endowment distribution source and liquidity uses? Does additional liquidity need to be maintained in operating reserves or in the endowment, and if liquidity is currently tight, could it be increased at a reasonable cost?

Especially if there is heavy dependence on the LTIP, in a stressed environment, what is the chance that operational shortfalls could require violating the spending rule to make extraordinary distributions? Or are there other firewalls in place to make that unlikely, such as stabilization reserves, the ability to cut spending, or to borrow?

Could debt covenants be tripped, or are key financial ratios on the edge of causing a credit downgrade and increasing future borrowing costs?

If the institution's revenues or expense exposures are highly concentrated, should the endowment avoid exposure to the same risks, or even hedge against them? Consideration of the rough correlations among the institution's revenue streams can be instructive. Taking the case of a university, is it likely that gifts and tuition would come under pressure during a downturn, creating heightened demands for endowment distributions when the LTIP is least able to weather increased outflows? (Answer: quite possibly!)

In addition to uncovering whether these institutional-level constraints exist for the LTIP, the Enterprise Review can also be useful in developing better context for forming a peer universe. Specifically, which institutions *should* be considered as peers for comparative purposes? Broad peer groups such as "all colleges and universities" are likely to encompass too much diversity to be truly useful, yet even more narrowly defined ones such as "all colleges and universities with assets between X and Y" will comprise institutions with quite different business models. But institutions with similar business models, comparable levels of endowment dependence, or operating exposures to a particular revenue stream would be closer peers. *Comparisons of asset allocations and investment returns for such peers would be far more meaningful than for "peers" defined mainly by endowment size.*

Apart from the hard facts, there is also the more qualitative question of how key stakeholders view these situations; that is, what is their risk tolerance? As risk is multi-faceted, this must be viewed through different lenses. How do they make the trade-off between portfolio volatility—resulting in variability in endowment spending, debt metrics, unrestricted monies, and any operating funds invested in (alongside) the LTIP—and the risk of failing to preserve long-term purchasing power? How comfortable are they with the constraints imposed by illiquidity in the LTIP? Most simply, which (if any) of these questions keep them up at night?

An Aerial View on Setting Investment Policy

As indicated earlier, traditional asset allocation fails to capture common risk factors. This can be rectified by placing a greater emphasis on top-down risk exposures.

The following summarizes the primary top-down characteristics included in investment policies under this framework.

ALLOCATIONS BASED ON ROLE-IN-PORTFOLIO EXPOSURE
Diversified Growth
Deflation Hedge
Inflation Sensitive
CONSTRAINTS AND OBJECTIVES
Illiquidity Tolerance
Currency Risk
Equity Beta
Value-Added Objectives

This example focuses on an endowment that has relatively high spending demands (e.g., roughly 4% or more of endowment market value adjusted for inflation). However, the framework is flexible and can be customized to apply to a variety of investor and pool types. For example, Appendix A discusses the way in which we frame a policy portfolio for pension funds. The primary objective is to understand how to manage key risks or liabilities and to consider such risk mitigation needs in parallel with desire to grow the LTIP.

Policy is based on desired role-in-portfolio exposures. These roles reflect our belief that a diversified exposure to equity-oriented investments is the best way to achieve long-term growth. However, such investments must be accompanied by dedicated protection against unexpected inflation and deflation/prolonged economic contraction to meet liquidity needs in turbulent periods when growth-oriented portfolios suffer.

Decisions regarding the specific ways in which the growth is achieved (e.g., asset allocation detail, manager selection) are reserved for implementation. *The policy outlines the most timeless and fundamental exposures—provided investors' circumstances don't change.* These exposures, when combined with key constraints and objectives, provide those implementing portfolios with the appropriate guidelines for meeting long-term goals.

EVALUATING TRADE-OFFS AND SETTING OBJECTIVES

When setting policy, you should begin with an understanding of the sorts of risks you are willing to take to achieve desired returns. *The framework looks sequentially at individual risk/return trade-offs to simplify and clarify the decision-making process. This is in contrast to typical asset allocation exercises that consider a complex set of trade-offs and implementation decisions together in a holistic manner.*

The framework looks sequentially at individual risk/return trade-offs to simplify and clarify the decision-making process.

The Enterprise Review facilitates informed decision making about the appropriate level and types of risk, risk preferences, and return objectives in the context of individual investors' circumstances. Most importantly, it should clarify the role of the LTIP. In general, an endowment serves one of three roles:

- maintenance of intergenerational equity;
- expansion of permanent capital; or
- capital spend-down to support strategic initiatives or to meet operating deficits.

The relationship between the spending rate and the portfolio's inflation-adjusted return objective is the most direct expression of that role. Maintenance of intergenerational equity requires that the portfolio earn a real return at least equal to its spending rate. Higher returns allow for expansion of permanent capital, as well as a margin of safety,[3] while returns that fall short of spending result in a capital spend down. The majority of endowments and foundations we work with seek the first objective—to maintain purchasing power while supporting current programs.

While return objectives relate directly to the role of the LTIP, volatility objectives are related to a variety of factors. They should be informed by an understanding of any constraints on the LTIP and your ability to absorb portfolio volatility. For example, how might short-term portfolio volatility affect the institution's level of spending? What is an appropriate trade-off between the risk of a real decline in short-term spending and the ability to preserve purchasing power over the long term?

Such analysis should be complemented by an evaluation of your *willingness* to incur volatility. Of course, most investors would want the lowest volatility achievable within the bounds of what is possible. However, the level of volatility taken to achieve return objectives has implications for how portfolios are implemented. These considerations should be explicitly taken into account as part of the policy setting process. The higher the return objective relative to portfolio volatility, the more diversification required to improve portfolio efficiency and the more sources of value added required to meet return goals.

[3] A "margin of safety" that is in place year after year but is never called upon, becomes, in effect, a policy of capital expansion.

Figure 5 **Considering Risk Trade-Offs**

Risk Trade-offs for "Simple" Stock/Bond Portfolios

Long-Term Risk: the probability of failing to earn a real return of 5% (policy spending rate) over 25 years.
Short-Term Risk: the probability of experiencing a 20% or greater cumulative decline at any point during five years.

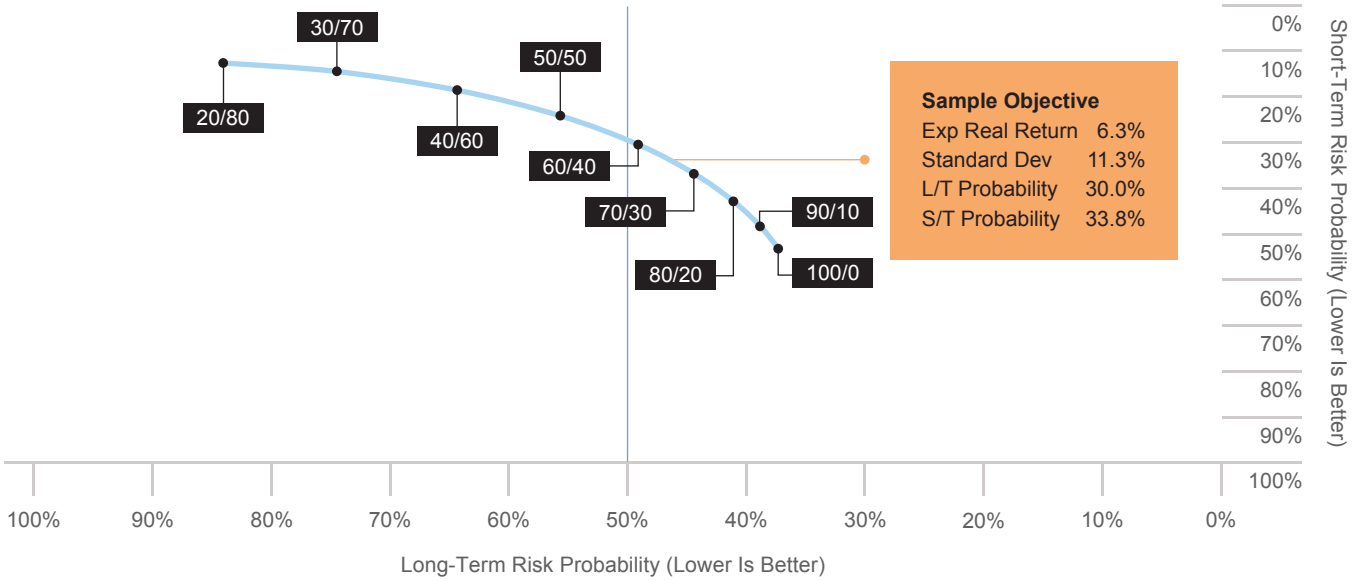

Notes: Based on long-term equilibrium risk, return, and correlation assumptions. "Stocks" represents 90% global equities and 10% emerging markets equities. "Bonds" represents 100% U.S. government bonds.

In evaluating investors' willingness to incur short-term volatility, we use a simple stock/bond portfolio to describe market volatility, called a "volatility-equivalent simple stock/bond portfolio."[4] Using such portfolios, you can look at the trade-offs between the risk of short-term declines in market value under stressed conditions[5] and the long-term risk of failing to preserve purchasing power to help clarify the degree to which you *value* reducing one risk over the other (see Figure 5).

In this context, if the return required exceeds what this volatility-equivalent simple stock/bond portfolio provides, it would have to be earned through value-added strategies in portfolio implementation. Of course, part of the policy setting process is also to evaluate if these goals are realistic given your attitudes about risk taking and available resources for portfolio implementation.

INCORPORATING PROTECTION AGAINST MACRO RISKS: BALLAST, NOT BULLETPROOFING

Equity-biased portfolios are clearly vulnerable to decline, particularly during macroeconomic shocks. Protecting against such shocks is costly, whether through allocations to assets that historically have proved defensive under such environments or through derivatives strategies that often come with negative carry.

Owning too little protection can be costly as well; therefore, you should size such positions in

[4] Note this is *not* intended to be an expression of a policy target stock/bond allocation.

[5] We use our stressed performance assumptions, which reflect more volatile conditions than do our equilibrium return assumptions.

direct relation to your risk tolerance and ability to navigate such difficult environments. You should carefully evaluate how much you are willing to commit to these sorts of protections. The answer may be, "not very much," which is actually common practice. Yet such allocations should be regarded as a way to fund liquidity needs—namely spending and capital calls (if any)—to reduce the need to sell risky assets at depressed prices during a crisis period. We believe that trying to fully protect the market value of portfolios from short-term declines is simply too costly for most institutions with a sufficiently long time horizon and high spending objectives.

In considering these defensive policy allocations, we define the positions in a relatively pure fashion. For example, high-quality sovereign bonds historically have been a reliable deflation hedge and have negative correlations to equities during economic contractions. Likewise, commodities and short-duration inflation-linked bonds also have low equity betas, and while less proven, have a reasonably strong relationship to inflation, provided valuations are not prohibitive.

Despite a history of providing defensive characteristics during periods of relevant macro stresses, these investments offer relatively limited value-added opportunities, and they have relatively low long-term expected returns. Currently, they are also quite expensive, raising questions about their ability to protect against macro risks from a starting point of today, but that is an issue we consider in implementation, not policy.

In contrast to policy, implementation decisions would include increasing or reducing deflation hedge allocations, implementing such allocations via shorter- or longer-duration sovereign bonds than reflected in the benchmark, or including credits. It also includes varying inflation-sensitive allocations and implementing via investments including significant interest rate or equity-like risk. Investors that choose to take such actions should be evaluated on their success or failure. Alternatively, investments such as opportunistic real estate or energy investments might be better classified as diversified growth, as that more accurately describes the role they fill in portfolios.

COMPLETING THE POLICY PICTURE

After volatility and return targets are set and the inflation-sensitive and deflation-hedge targets are established, we stress test for the ability to take on illiquidity and currency risk. These fundamental objectives and constraints are codified as investment policy. The framework uses a simple, investable policy benchmark to capture the following three role-in-portfolio exposure allocations without any direct expression of more active implementation decisions:

- diversified growth-oriented investments, benchmarked to equities with a beta adjustment;
- inflation-sensitive assets, benchmarked to short-duration inflation-linked bonds and commodities; and
- deflation-hedge assets, benchmarked to intermediate- to long-duration high-quality sovereign bonds.

We include an equity allocation with a beta of less than one to incorporate a key implementation objective—reduce equity beta risk both through diversification and through pursuit of strategies that add alpha.

Implementation—It's All Relative

The implementation stage of the Risk Allocation Framework is distinct from the policy setting phase. Implementation includes all decisions that result in differences between the actual portfolio and the policy portfolio, including more detailed asset allocation, manager structure, and manager selection. Policy is intended to be timeless, while implementation is more dynamic.

Implementation should start with a simple question: *What positions will I take to add value relative to the policy portfolio?* Such positioning should reflect a myriad of factors including, but not limited to:

- competitive advantages (e.g., ability to identify and access managers with high alpha potential);
- views on strategic investment themes (e.g., overweight emerging markets within global equities);
- relative valuation differentials;
- practical constraints (e.g., availability of open or fund-raising high-quality managers in areas of interest); and
- the desire to further diversify exposures.

The treatment of asset allocation detail as implementation, not policy, elevates the importance of maintaining top-down exposures that are more essential to meeting investment objectives, while clarifying the intent to be more flexible and dynamic on asset allocation defined as implementation. *In fact, as equity betas and volatility levels change meaningfully, you should alter exposures to asset classes or managers to bring these risk characteristics in line with portfolio objectives.* This should be considered only at extremes and not in isolation. For example, a higher than normal equity beta may be desirable when equity valuations appear to be at very depressed levels, while a below-target equity beta would be desireable at extreme valuations.

Likewise, different asset classes with common risk factors may be more or less attractively valued over time. Treating asset allocation as implementation improves your ability to invest in what is most attractive without the requirement to maintain allocations to asset classes that may share essentially the same risk factors and economic drivers of returns.

Policy is intended to be timeless, while implementation is more dynamic.

For example, if you are interested in distressed real estate opportunities, this framework encourages comparison of opportunities across "asset classes" such as public long-only managers, hedge funds, or opportunistic private real estate investments. This is particularly important for relatively illiquid assets that require a long-term commitment. On occasions when equity beta is cheap and/or comparable private strategies are particularly expensive, you should think carefully about taking the incremental risk of long-lockup private investments.

Maintaining asset allocation implementation targets (and rebalancing to them) is still a useful part of portfolio management. The objective is to distinguish sharply between exposures that are critical to maintain and those that can be allowed to vary when implementing top-down portfolio objectives. A commitment to review asset allocation at least annually and to think about equity and equity-like exposures together allows for a more comprehensive exploration of the best ways to implement.

Figure 6 **Where Are You Taking Active Risk?**

Manager-Weighted Active Risk of a Sample Portfolio (Ranked by Current Active Risk)

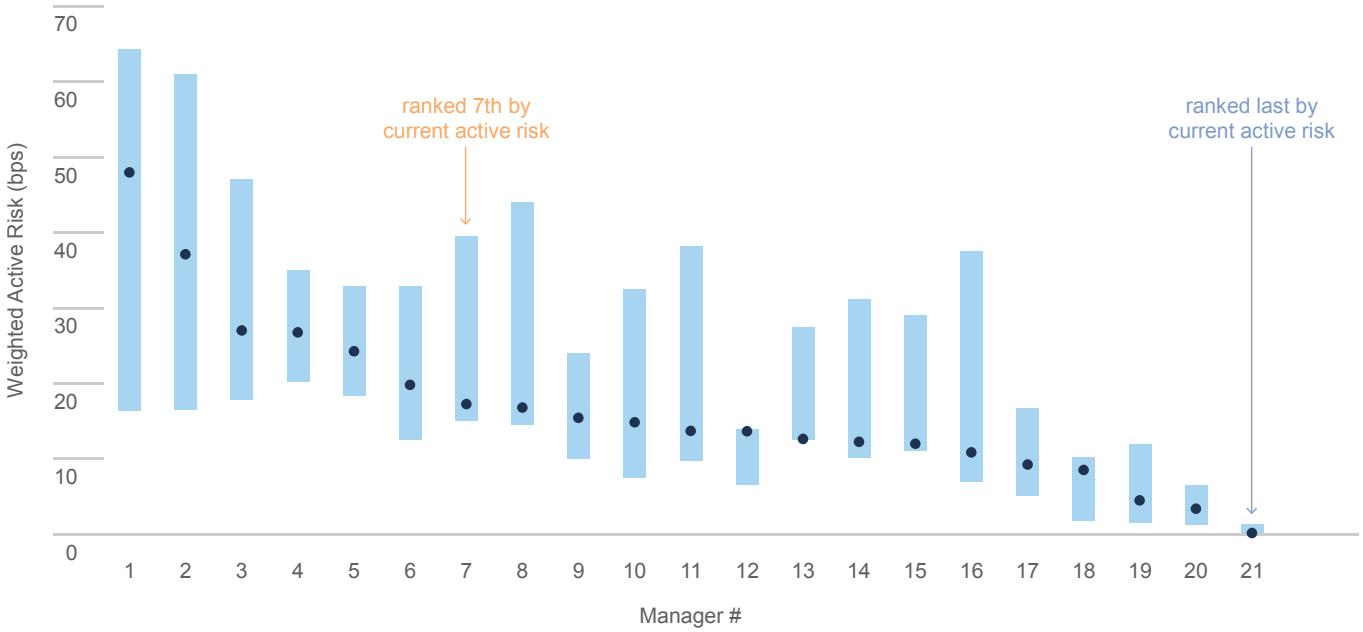

Manager Weight in Sample Portfolio

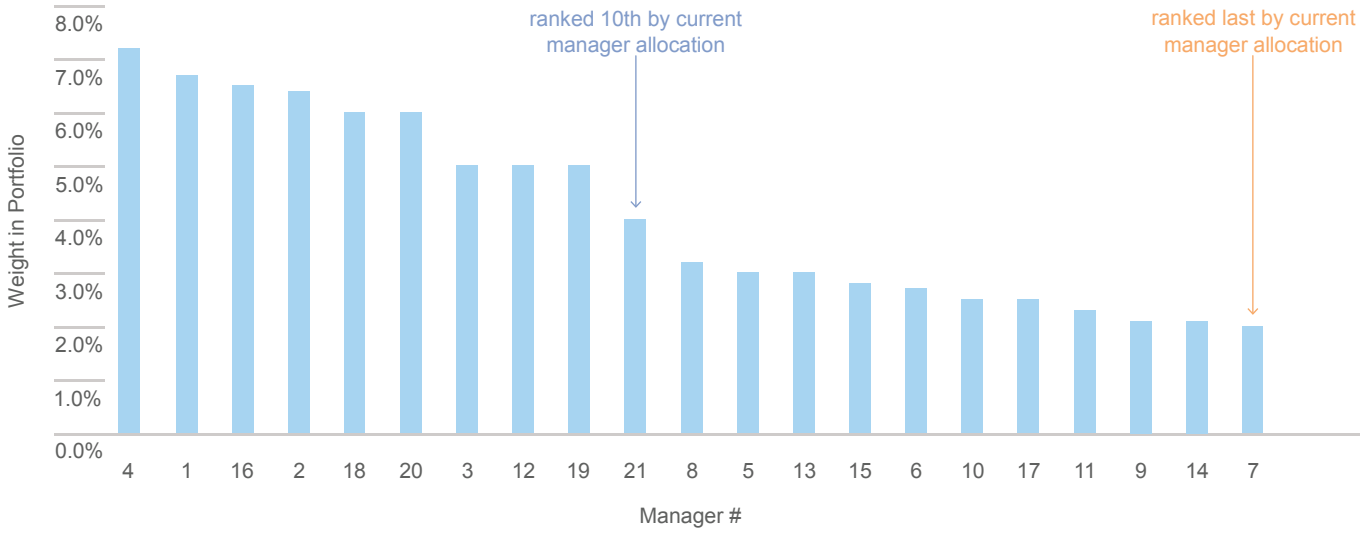

Notes: For the top chart, the floating bars represent historical min/max, and the dark blue circles are the current environment active risk measure. Active risk and weights in portfolio reflect liquid and semi-liquid investments only; weights in portfolio will not sum to 100%. The range over which min/max is determined may vary across managers. Analysis is based on rolling 24-month exponentially weighted statistics.

A DYNAMIC APPROACH TO MONITORING IMPLEMENTATION RISK

All of these activities are employed with the goal of adding value over policy. However, these differences also introduce active risk. The higher the active risk, the greater the opportunity to make an impact on portfolio performance. However, active risk can cut both ways and should be expected to produce underperformance relative to policy at times.

Understanding at each decision level (e.g., asset allocation, manager structure, manager selection) how much active risk you are taking, how your decisions affect the total portfolio equity beta exposure, and how they affect total portfolio volatility is an important part of portfolio construction best practices.

The portfolio should be monitored on an ongoing basis to make sure risks taken are consistent with convictions about their potential value added relative to policy. For example, Figure 6 shows the active risk and the capital allocations of managers in a portfolio. Capital allocations alone do not provide a full picture of the potential impact a manager (or other position) can have on the portfolio. Note how the ranking of managers varies when ranked by capital allocation or active risk. For example, the smallest manager by weight in the portfolio is in the top third of managers in terms of active risk.

Given the changing nature of implementation risks and potential rewards, the Risk Allocation Framework incorporates dynamic risk monitoring tools on the liquid and semi-liquid part of the portfolio that can be managed over shorter-term horizons. (Please see Appendix B for a brief discussion of treatment of private investments.) These tools allow an evaluation of current environment risk exposures, such as standard deviations, active risk, and equity betas for the total portfolio and each position in the portfolio down to the manager level. We weight data exponentially such that recent periods have greater influence than older ones to capture current changes in markets more than do traditional calculations that equally weight historical observations (see Figure 7).

Such quantitative analytics provide an additional lens through which to evaluate the portfolio implementation decisions. They should be used to complement qualitative knowledge about markets and managers, market cycles, and other risk factors, such as valuations. Such risk metrics should be considered in the context of rolling historical data to understand (for example) if an asset class or manager's low equity beta is typical or at extremes relative to historical experience.

Understanding current relationships among asset classes and managers can provide insights about risks taken today.

Figure 8 provides an example of rolling active risk and equity beta statistics for a total portfolio. In this stylized example, all asset class and manager allocations remained constant, yet active risk increased meaningfully in late 2008, gradually returning to pre–global financial crisis levels. Such changes in active risk levels should prompt you to evaluate the underlying managers to understand the source of the elevated active risk and whether it is desirable in the context of the current environment.

While it is useful to incorporate static assumptions about long-term returns, correlations, and standard deviations, particularly in policy setting intended to represent long-term risks and objectives, we know that such characteristics are unstable over shorter time periods. Correlations among all equity asset classes have increased significantly

Figure 7 **Using Exponential Weighting to Enhance Understanding of the Current Environment**

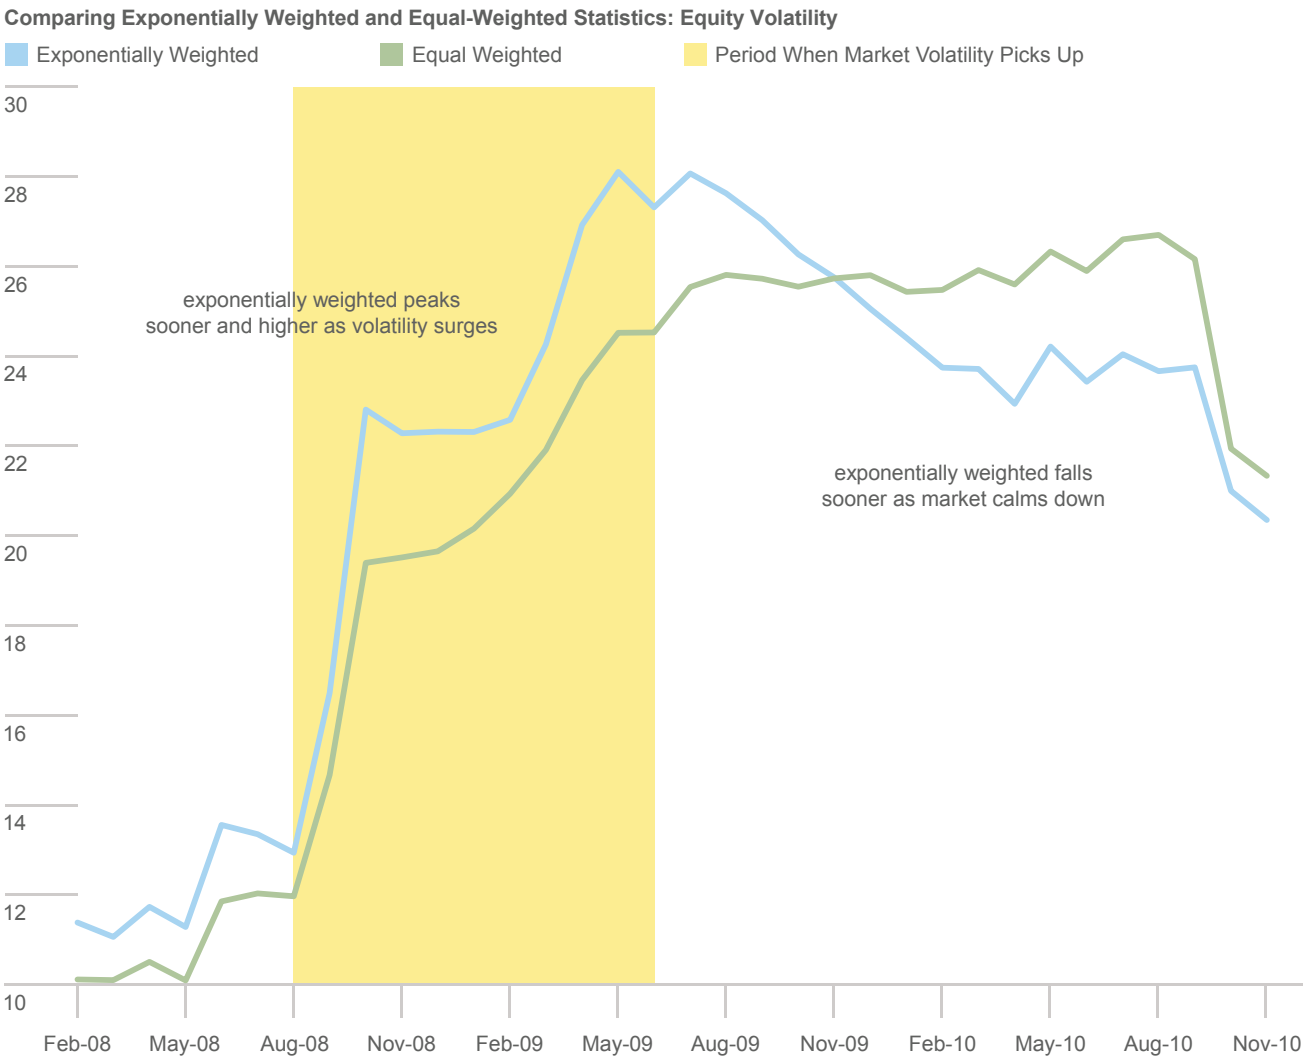

Exponential weighting can be used on historical data streams to better capture current risk by placing more emphasis on the most recent data.

Our methodology applies a 12-month half-life exponential weighting curve to a window of 24 monthly data points, meaning that a 12-month-old data point is weighted half of what it would be as the most recent month. This differs from conventional volatility and regression statistics, which use all data points equally, so that a data point has the same impact on the calculation no matter where it falls in the 24-month window. We use a 24-month window as it passes basic data sufficiency requirements, but is short enough to capture new developments in the environment.

As the graph highlights, exponential weighting helps statistics pick up on changing trends faster, both in reflecting surges in volatility and moving on from them, and is therefore better aligned with the goal of capturing risk in the current environment. The chart shows rolling historical equity volatility and demonstrates how the increase and subsequent decrease in volatility is reflected sooner in the exponentially weighted statistic than in the conventional, equally weighted statistic.

over the last several years relative to historical averages, revealing that models based on long-term assumptions overstate diversification benefits at present. *Understanding current relationships among asset classes and managers can provide insights about risks taken today.* For example, when correlations are elevated, portfolio volatilities rise relative to historical levels. This suggests that lowering equity betas could be beneficial in bringing total portfolio volatility down to desired levels. Such analytics also help identify changes in manager risk levels more quickly.

Recognizing that not every position taken will pay off all the time, and in fact some of the bets may be negatively correlated with each other, also provides useful insights into what to expect in terms of relative performance. Of course, risks come with downside and upside, so it is essential to be selective in your implementation decisions and to seek to understand and diversify the risks taken to meet objectives.

You can improve your likelihood of meeting long-term investment objectives by paying close attention to the risk/reward trade-offs of investments, understanding how they should perform in different environments, and recognizing that downside risk is inevitable. Such analysis can also help you stick to your strategy when times get tough.

Figure 8 **Monitoring Implementation Requires a Dynamic Approach**

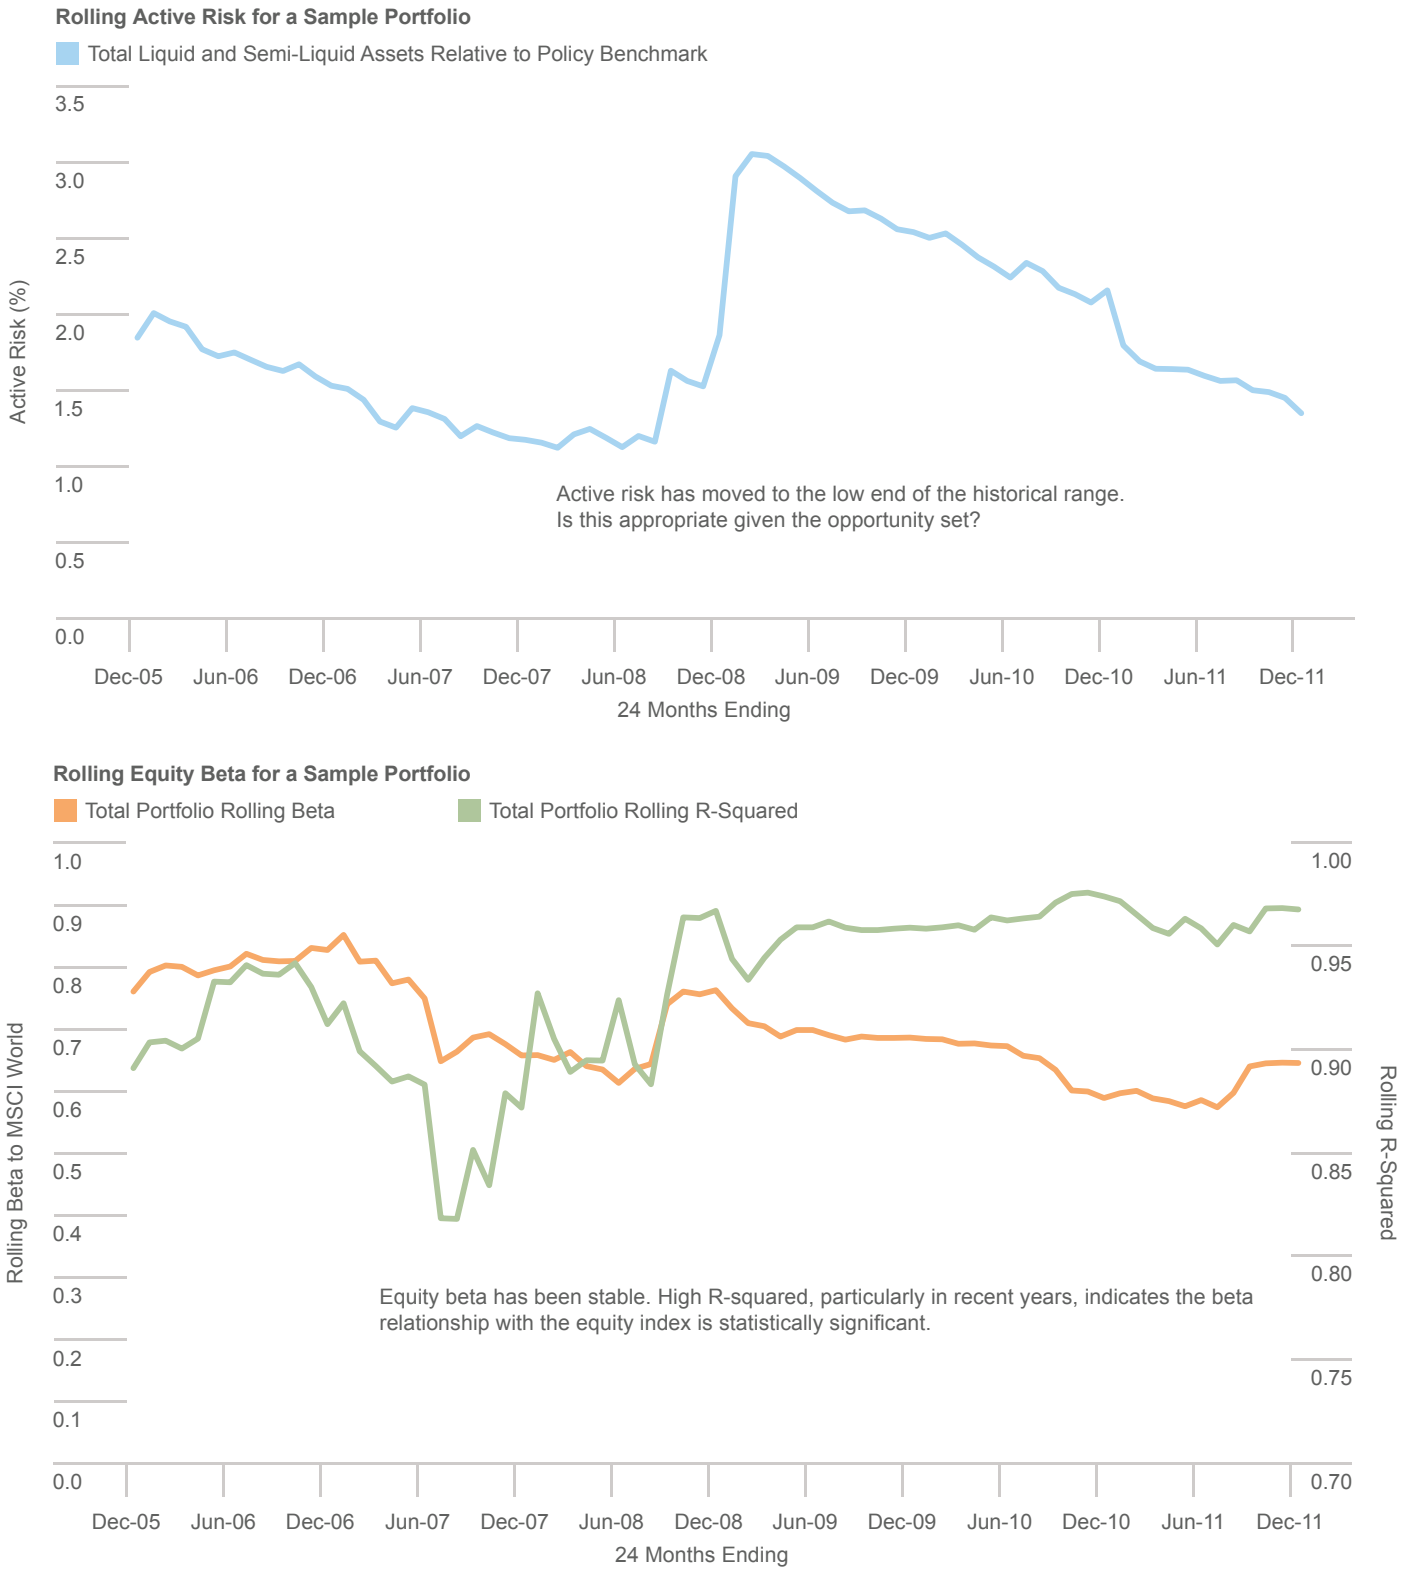

Notes: Rolling statistics in this exhibit are all exponentially weighted. Please see Appendix B for a discussion of the treatment of private investments.

Measuring Performance to Create a Virtuous Feedback Loop

The Risk Allocation Framework uses ongoing performance monitoring and attribution to shed light on the ways in which value has been added to (or subtracted from) the portfolio. Repeated reviews of value added should help to provide clarity on what sorts of practices should be embraced, improved upon, or avoided.

The most fundamental measure of performance is evaluating the ability to meet long-term objectives. This metric should be considered over no shorter than rolling five-year periods—the longer, the better. We make this assessment by comparing the performance of the portfolio with a benchmark of its volatility-equivalent simple stock/bond portfolio. Evaluating portfolio value added in these terms measures whether you are meeting long-term goals, given what the capital markets provided. In contrast, comparing portfolio performance to an absolute-return type of benchmark, like a 5% real return, creates many difficulties, as it presumes no market volatility (see next page).

Repeated reviews of value added should help to provide clarity on what sorts of practices should be embraced, improved upon, or avoided.

Similarly, we look at performance of the policy benchmark versus its volatility-equivalent simple stock/bond portfolio benchmark over relatively long time horizons to evaluate the value added from having that policy.

Finally, we look at various implementation decisions relative to policy. We can consider these in aggregate by comparing total portfolio performance to the policy benchmark, and we can attribute value added to asset allocation and manager selection. The levels you look at should be consistent with your decision-making process. For example, if you don't make discrete manager structure decisions, but rather consider that as part of your asset allocation process, you might combine those two levels of value added together in your evaluation process. By creating benchmarks to evaluate each level of portfolio decision making we can look at their contributions to performance independently.

Performance monitoring should serve as a useful feedback loop to the decisions made at the policy and implementation levels. You should seek to learn from evaluating what you have done well (and not so well), to understand what can be improved upon, and to incorporate such lessons into your risk allocation practices.

THE CHALLENGE OF EVALUATING A FIXED-RETURN OBJECTIVE (E.G., CPI + 5%)

Evaluation of performance relative to policy objectives is a challenging exercise. Such assessments require long time horizons and typically use benchmarks that are often ill-suited to the task. Benchmarking is most effective when the performance being evaluated is relative to a specific, well-defined mandate that is both objectively representative of the opportunity set and passively investable.

- Portfolio long-term return objectives are typically expressed as fixed-return objectives (e.g., CPI + 5%). Such fixed-return objectives do not represent an investable opportunity set. As a result, they are ineffective tools for evaluation—at least over time horizons that are relevant for oversight purposes.
- A fixed-return objective does not consider the general investment environment—e.g., during difficult market periods, there may not be any reasonable policy portfolios that could achieve the return objective. Conversely, during periods of strong investment returns in capital markets, "merely" achieving the return objective would be a disappointing result.

SIMPLE VOLATILITY-EQUIVALENT BENCHMARKS HELP ADDRESS THESE CHALLENGES

- The range of simple equity/bond benchmarks can serve as a scale for characterizing a portfolio's volatility profile. In this case, it would explicitly not be a reflection of the actual portfolio's split between "equity-like" and "bond-like" investments and would be independent of the actual decision-making process.
- The value added from the decision to adopt a certain policy is then defined as the performance of the policy benchmark relative to its volatility-equivalent simple stock/bond portfolio.
- The ability to evaluate performance over shorter time horizons is improved since the simple portfolio represented in the benchmark reflects the return available in the current environment given a very basic implementation of a desired level of portfolio volatility.
- Merely meeting the performance of a volatility-equivalent simple stock/bond portfolio should not be considered success. The objective of policy setting and implementation is to add return over what you could earn by simply investing in a stock/bond combination with the same volatility profile.

Governance Matters

While not new or unique to our Risk Allocation Framework, we emphasize that without appropriately strong governance, the best constructed investment plans could easily fall short of the mark. Both at the institutional level and specifically pertaining to the LTIP, well-crafted governance increases the likelihood (but alas does not guarantee) that decision making will be sound, and—importantly during market stress periods—that it will be timely, coordinated within an institution, and resilient against the behavioral risks to which investors are subject in crises.

As we explored in our 2009 paper, *Behavioral Risk*, investors are not necessarily the rational actors of economic theory, but human beings, with hard-wired emotional and psychological reactions to shocks. Such responses include increased risk aversion, a desire for liquidity, a shortening of the investment time horizon (including extrapolation of current trends and sudden amnesia about the lessons of market history), and heightened unwillingness to be contrarian. Recognizing that this happens, what can be done about it? The mistakes stemming from behavioral risk can be mitigated through governance that simplifies decision making and that instills self-awareness and discipline, through education about expectations and long-term market history, and through well-crafted portfolio design that focuses on the different types of risks taken to achieve return objectives, and their trade-offs.

Best practice governance helps to instill discipline by setting expectations of what could happen well in advance so that decision makers are prepared to act. In other words, by making it less likely that decision makers will be surprised and overreact, governance can enhance the resilience of the LTIP. This applies not only to decision makers, but also to external stakeholders. The most capable decision makers may find it challenging to take the right action while faced with harsh stakeholder opposition. Management of such risks is perhaps best achieved via continual education in the lessons of market history as they pertain to the policy portfolio, which should be reviewed on a regular schedule, ideally annually.

Points to be reiterated include: the intended roles of the investments; the long-term return and volatility expectations for the portfolio and the likelihood of achieving objectives; and the expected frequency, severity, and duration of declines. Re-affirming all of these should help make it easier for stakeholders to remember why the policy was developed in the first place, and may lend it greater weight and enable it to have more staying power.

The mistakes stemming from behavioral risk can be mitigated through governance that simplifies decision making and that instills self-awareness and discipline.

In addition to educating decision makers on how they *should prepare* to act, the best governance also serves to remind decision makers whether or not they *ought* to act. Detailed minutes should document what decisions were made and the reasons behind them, and what decisions were considered but not acted upon, and should be paired with performance reports for evaluative purposes. Continual assessment of value added by decision makers should help to remind them of where their comparative advantages lie, and what sorts of decisions they should avoid making, especially during periods of high behavioral risk.

Other important governance issues that bear upon risk management include board, committee, and organizational structure; communication among key decision-makers; management and investment of operational liquidity; and enterprise strategic planning.

Evolution or Better Packaging?

Skeptics might question whether the Risk Allocation Framework is an evolutionary step forward or merely just the same old fellow, dressed in a new suit. Asset allocation—informed by rigorous valuation analysis—and manager selection remain an important part of the Risk Allocation Framework, so there is some family resemblance. In fact, this is by design, as we preserved the best practices of traditional portfolio construction (often referred to as "the endowment model" because endowments were the earliest adopters).

However, the Risk Allocation Framework's primary merits are that it enables investors first to more rigorously *construct* portfolios designed to realize their investment objectives and second to understand much more clearly how best to *manage* such a portfolio dynamically to improve the likelihood that it performs as anticipated.

Each step of the framework builds on those that come before:

- The systematic approach of the Enterprise Review helps you to understand the linkages between the LTIP and the institution to improve the LTIP's ability to meet institutional demands during normal and stressed periods.
- The insights gained from such reviews inform policy setting and facilitate your ability to set return objectives and constraints thoughtfully to suit the institution's needs and risk attitudes.
- The guidelines set in the policy provide limits on exposures to common risk factors that cut across asset classes. This provides the implementation flexibility you need to invest in what is currently most attractive, while providing a clear understanding of constraints at the most fundamental level. Given that such risk factors shift over shorter time horizons, you must supplement this flexibility with ongoing and rigorous monitoring, and make adjustments as appropriate.
- Ongoing performance measurement and attribution analysis complete the circle. You should study the success and failure of decisions to learn from the past and seek continuous improvement.

These enhancements to the portfolio construction process clarify your understanding of portfolio exposures in a complex world of proliferating asset classes with varied names, but often similar economic bases of returns. The monitoring process can be time consuming, but the benefits you gain from an improved understanding of the nature of risks taken in the portfolio and how they relate to each other and to the investment policy are worth the effort. ■

Appendix A Applying the Framework Beyond Endowments

This paper uses endowment examples to introduce the Risk Allocation Framework. However, the framework is flexible and can be customized to apply to a variety of investor and pool types. For example, our portfolio construction process for defined benefit pension plans under a Liability Driven Investment (LDI) framework is consistent with the Risk Allocation Framework. The primary differences with the process for endowments relate to the nature of issues investigated in the Enterprise Review and the definition of the objectives and constraints (specifically as they relate to liability-related risks and time horizons), as well as role-in-portfolio allocations specified in the policy setting process. We provide an overview of these steps in the process as they relate to defined benefit pension plans to give a sense of the adaptability of the Risk Allocation Framework.[1]

WHAT IS LDI?

An effective LDI framework allows an institution to evaluate asset allocations and portfolio implementation in the context of its relevant liability and unique organizational circumstances and risk tolerance. The framework seeks primarily to generate portfolio returns sufficient to fund the contractual liability and, in most cases, generate some excess return, but to do so in a risk-controlled manner. Theoretically, this framework will result in a more efficient investment solution and superior risk management, allowing institutions to better balance the potential rewards of higher returns with investment and organizational risks.

LDI approaches focus on managing the relationship between the size of the asset pool and the related liability. This relationship is often referred to as surplus and the volatility of this relationship is referred to as surplus risk. The foundation of an LDI approach is assessing the sensitivities of the assets and liabilities to a variety of factors such as changes in interest rates, inflation, and a broad range of capital market environments. It also considers an institution's financial health and the economic sensitivity of an institution's operations or business.

ENTERPRISE REVIEW

For defined benefit pension funds, the Enterprise Review would explore the plan sponsor's return objective and explicitly identify its ability and willingness to assume risk (risk tolerance). Sponsor risk tolerance, return goals, liability characteristics, and plan terms are the key parameters around which an LDI policy, or risk-budgeting framework, is designed.

For defined benefit plan sponsors, the concept of risk tolerance generally focuses on identifying an acceptable range of funded status volatility, which derives from institutional tolerance for contribution and balance sheet volatility. An acceptable level of volatility can be defined either in percentage or monetary terms. Within an LDI process, funded status volatility is often referred to as surplus risk. A simple way to frame the issue of risk tolerance is to determine whether to make contributions more regularly but with a smaller range of potential contributions, or to make larger contributions in the hope of making them less frequently.[2] If a plan sponsor is more comfortable with a smaller range of potential contributions, then it would operate with a lower risk budget. However, sponsors would seek to maximize returns within the parameters of the relevant risk tolerance.

For a sponsor to understand its ability to take risk, it needs to understand how the plan's surplus risk interacts with its operating risk. In most cases, this entails understanding the relationship between the sponsor's operations, capital market returns, and changes in market-based interest rates. Therefore, it is helpful to keep in mind the circumstances under which defined benefit pension plans experience the largest negative tail events—deflationary recessions

[1] Readers interested in learning more about our portfolio construction process for pensions should review our 2011 report *Pension Risk Management.*

[2] In the latter scenario, these contributions are more likely to take place at peak periods of stress for the plan sponsor's operating business.

or depressions when interest rates or discount rates decline significantly, and risk assets generally perform poorly. A number of plan characteristics can lower a sponsor's risk tolerance (and the more of these characteristics that an institution has, the lower the risk tolerance is likely to be). Considerations to evaluate include:

- size of plan liability relative to the size of sponsor's balance sheet;
- potential size of future contributions relative to the organization's projected free cash flow;
- correlation of operations to the return of risk assets and changes in interest rates;
- correlation of potential lump sum payments to drawdowns in plan surplus and sponsor financial health; and
- funding time horizon of plan.

A sponsor's risk tolerance is also defined by institutional *willingness* to assume risk and derives from the psychological and behavioral dynamics of the sponsor and the individuals responsible for plan oversight. Risk tolerance can be quantified simply as the monetary amount of surplus volatility an institution is willing and able to assume, with the constraint being the lesser of the ability or willingness to assume risk. This includes evaluating tolerance over various time horizons (annually, rolling three years, etc.) or focusing on tail event surplus risk.

POLICY SETTING

After defining risk tolerance, the focus shifts to connecting the assets and liabilities within the investment process. There are significant uncertainties associated with the future value of both the plan's assets and the plan's liability. Over time, changes in the value of the liability, excluding the effects of future accrued benefits, will be driven in large part by changes in interest rates, inflation, and variability around mortality assumptions. Investment decisions should be made using a comprehensive risk framework that evaluates asset returns and volatility relative to changes in a plan's liability.

To create a framework that allows for an evaluation of risk relative to a liability, sponsors must understand the "risk-free," or risk-neutral, position. The simplest way to define the risk-free asset pool is generally to identify the theoretical asset pool that "perfectly" hedges the liability. This theoretical risk-free asset is the zero relative volatility asset for an investment portfolio, which differs from a purely asset-based perspective where the theoretical risk-free asset is often considered high-quality sovereign cash. By using the liability, or risk-free asset, as a benchmark for plan assets, investment decisions become a risk-budgeting process that evaluates the trade-off between expected return and risk relative to the risk-free asset pool.

In very simple terms, forming an asset allocation policy for a defined benefit plan involves a risk-budgeting process that weighs the decision of allocating assets between the following two theoretical portfolios:

- *Hedging Portfolio.* This portfolio attempts to minimize surplus risk.
- *Growth Portfolio.* This portfolio attempts to generate excess returns that sponsors target to reduce contributions. In essence, the excess return is used to offset a portion of future accrued benefits, thus reducing sponsor contributions. A significant portion of a plan's risk-budgeting process will focus on creating a diversified growth portfolio of beta and active risk exposures. Importantly, the active risk exposures should be relatively uncorrelated to capital market betas and changes in interest rates. These active exposures include tactical asset allocation, manager selection, and manager structure.

A holistic pension risk-budgeting process should also focus on the levers within the two portfolios to create a capital- and surplus risk–efficient portfolio. For

instance, simpler LDI frameworks assume that all growth portfolios have the same characteristics (e.g., exposures, excess return, and risk), which is obviously not necessarily true. By using levers within the growth portfolio, such as diversifying across beta and alpha sources and allocating more or less risk to beta and active components, growth portfolio surplus risk can be altered, thus changing the size of the hedging portfolio required to obtain a given level of risk. This allows for an array of plan portfolios that look distinctly different, but have similar expected liability relative to risk profiles. In the end, this process should focus on maximizing expected excess return for portfolios based on a sponsor's acceptable level of surplus and institutional tail risk.

Assuming institutions believe that they, or their advisors, have the skill and resources to identify active strategies and managers that add value, sponsors should create targets for various beta exposures and for active risk exposures. Successfully allocating additional risk to active sources of return or to betas that are less correlated to changes in interest rates allows for the creation of a more efficient surplus risk/return portfolio. Creating portfolios with a more efficient surplus risk/return profile allows sponsors to increase the size of the growth portfolio and decrease the size of the hedging portfolio, thus increasing returns at a given level of surplus risk. Importantly, this more efficient portfolio also allows sponsors to maintain the same allocation to their growth and hedging portfolios, thus maintaining a similar level of expected return at a lower level of surplus risk. We would emphasize that a plan oversight strategy that allocates a significant amount of risk to active exposures must extensively diversify sources of active risk, or total plan surplus risk may increase.

SAME PHILOSOPHY, DIFFERENT SPECIFICS

Once investment policy is set, implementation can proceed. As is the case for endowments, ongoing monitoring and performance measurement and attribution are critical parts of the framework for defined benefit pensions.

We discussed our LDI philosophy and its fit into the Risk Allocation Framework with regard to defined benefit pensions. However, the broad framework and strategies discussed can be adapted to effectively address asset-liability management for other retirement plans and other asset pools that support contractual liabilities. Further, the Risk Allocation Framework can be customized to apply to virtually any other type of investor with the goal of increasing the linkages between investment policy and your fundamental return objectives and risk tolerances while allowing for dynamic implementation accompanied by rigorous ongoing monitoring. ■

Appendix B Private Investments: Equity Betas, Active Risk, and Value Added

Ongoing dynamic monitoring of equity betas and active risk of private investments is of limited use, as such investment allocations cannot be explicitly managed beyond sizing capital commitments (or use of secondary markets). Further, quarterly returns—the standard in the private investment industry—that are made available on a lagged basis and are not regularly marked to market provide limited information about volatility and equity beta, making ongoing monitoring impractical. Given these limitations, we take a long-term perspective in assessing the contribution of private investments to equity betas, active risk, and value added relative to public market indices.

With regard to equity betas, we make a simplifying assumption that private investment equity betas can be proxied by those of marketable indices. We use broad indices given that the imprecision of this exercise precludes use of narrowly defined marketable proxies. For example, we would proxy most growth-oriented private investments (e.g., venture capital, buyouts, growth equity, opportunistic real estate, and energy investments) as global equities (e.g., using the MSCI World or ACWI equity index) and expect they would have an equity beta of 1.0. While you can argue that equity betas should be higher to reflect a higher degree of leverage, operating risk, etc., we think an equity beta of 1.0 is a reasonable assumption, particularly given that such measures are not observable on an *ex post* basis. Strategies that intuitively have less equity beta, as their economic basis of return is more closely tied to hard asset exposure (e.g., upstream oil & gas, core real estate), could be proxied by a marketable index associated with the appropriate sector.

To account for the degree of active risk inherent in private investments, you should make some assumptions based on the active risk experience of the private investments included in the portfolio (to the degree it is reasonably reflective of private investment exposures looking forward). This active risk level should then be deducted from the total active risk you intend to take so that you can consider the "risk budget" available for the marketable, more liquid portion of the portfolio.

Finally, given we assume that all equity-oriented private investments have an equity beta of 1.0, any value added over public market equities (measured over an adequately long time horizon of five years or longer) is presumed to be primarily from alpha sources. While some of this is likely strategy betas (i.e., the return reflective of the strategy as implemented by the "average" manager) and may also be market beta, we consider this a reasonable framework for measuring value added over the long term, particularly since public market equities tend to be the investment of choice if suitable private investments are not available. A similar case can be made for comparison of private investments in hard asset–related strategies that are benchmarked to commodity, natural resources equity, or REIT indices. ■

Appendix C Beware of Model Abuse

In recent years, we have developed many models to facilitate our portfolio construction efforts. While models are extremely helpful in informing decision making, it is important to understand their limitations. Models are expressions of theories, which in turn are generalizations (based on observations and experience) that attempt to represent "reality." Models therefore are a few levels of abstraction removed from reality and the two should not be confused.

We have long held the view that models of capital market behaviors are limited in their applicability. They test conclusions arrived at by other means, offer directional and rough-magnitude insights into the trade-offs among alternatives, and provide a useful framework for discussion and analysis; they do not provide answers. Their limitations derive from a number of factors. One is that not all inputs can be quantified precisely enough to justify their incorporation into a model, or if they must be incorporated, the model's outputs become correspondingly less precise. We acknowledge this in our advocacy of a fuzzy "efficient region" rather than a bright-line "efficient frontier" of best-possible risk/return trade-offs—we view the precision of the underlying inputs (asset class returns, standard deviations, and cross-correlations) as being sufficiently uncertain that the fairest portrayal of the outcome must also demonstrate some uncertainty.

Another weakness is that models are expected to have a "sweet spot"—which may indeed be a wide range—where they are most accurate, but will not perform equally well in all circumstances. Several of Cambridge Associates' models are mean-variance models (with lognormal distributions), which are effective at portraying the probability of results likely to occur most of the time, but they do not adequately portray "fat-tail" events. These are extreme events that occur more often than would be predicted by a lognormal probability distribution (or equivalently, where the outcome is more severe for a given likelihood than would otherwise be expected).

We do not attempt to adjust these models to better portray extreme outcomes, as this generally involves making additional assumptions and incorporating more complexity with no guarantee of more "accurate" results. When considering tail risk or extreme returns, we believe the probability of a bad event occurring is far less important than how such an event might affect the portfolio. We have therefore developed and use other tools not based on mean-variance to test scenarios and stress portfolios in extreme cases.

A third drawback of models is that, based as they are on observations of past experience, they cannot by their very nature account for never-before-seen events. Capital markets models based on "ordinary times" perform less well in the face of disasters and other historical discontinuities, including the unforeseen outcomes of financial innovations.

In short, we can use models and other analytical tools to measure risk and to compare the relative merits of different approaches. But to us, the above only strengthens our stance that models are a secondary complement to, and cannot substitute for, investment choices based on well-grounded qualitative theory. Reliance on models and extrapolation of the past into the future is not risk management. Understanding the limitations of past patterns of historical performance, using one's imagination to conceive of how the world may change, and taking action accordingly—that is risk management. ■

Appendix D Glossary

ACTIVE RISK

Active risk is the amount of risk taken relative to policy to permit value-added returns. It can be used as an *ex ante* (before the fact) measurement of potential impact of a position, or as an *ex post* measure.

Active risk is often referred to as *tracking error*. The risk is statistically defined as the standard deviation of value-added returns (portfolio return minus benchmark return). This can be calculated for the total portfolio or any individual active position (e.g., manager performance in relation to manager benchmarks). It is used to measure the degree of consistency of value-added sources and the potential impact that value added may have on the portfolio. A reference to the *amount* of value added should always be included for a meaningful interpretation and discussion of this metric.

ATTRIBUTION ANALYSIS

Performance attribution decomposes portfolio returns into different sources (e.g., asset allocation, manager structure, active management). Multi-period attribution is used to determine the amount of return from these sources over multiple periods.

ENDOWMENT MODEL

An approach to investing that adheres to the following key principles:

- long-term investment horizon;
- asset allocation policy consistent with both financial objectives and resources available for effective implementation;
- high allocation to equity assets to support spending needs without depleting the ability to meet the needs of future generations;
- precautionary hedges against "fat-tail" macroeconomic risks;
- adherence to value-based investment principles; and
- rigorous and ongoing monitoring of investment results.

ENTERPRISE

The overarching entity within which a long-term investment portfolio (LTIP) exists. As part of the Enterprise Review, we seek to understand the role of the LTIP within the business model of this larger enterprise (or business) and any constraints that it may impose on the LTIP. Note: the term "enterprise" is interchangeable with "business" depending on the nature of the entity under review. In the context of families, the "enterprise" would refer to the multiple branches of families in relation to their respective balance sheet, income statement, and cash flow needs.

EQUITIES WITH A BETA ADJUSTMENT

Adjustment made to equities (with a beta of 1.0) to represent some level of de-risking that comes from portfolio implementation through diversification.

From a policy benchmarking perspective, the diversified growth role category may be represented as global equities with a beta adjustment of (for example) 0.8, to reflect the addition of diversifying investments beyond pure equities. This would be benchmarked with an 80% weight to a global equity benchmark and a 20% weight to a cash benchmark.

EXPONENTIAL WEIGHTING

Exponential weighting is a methodology that better captures current risks by placing more emphasis on the most recent data. Within our analysis this methodology applies a 12-month half-life exponential weighting curve to a window of 24 monthly data points, meaning that a 12-month-old data point is given half the weight of that of the most recent month. This differs from conventional volatility and regression statistics which use all data points equally, so that a data point has the same impact on the calculation no matter where it falls in the 24-month window. A 24-month window is used

as it passes basic data sufficiency requirements, but in short enough to capture new developments in the economic environment.

IMPLEMENTATION

The implementation stage of the Risk Allocation Framework is distinct from the policy setting phase. Implementation includes all decisions that result in differences between the actual portfolio and the policy portfolio, including more detailed asset allocation, manager structure, and manager selection. Policy is intended to be timeless, while implementation is more dynamic.

Implementation should start with a simple question: *What positions will I take to add value relative to the policy portfolio?* Such positioning should reflect a myriad of factors including, but not limited to:

- investors' competitive advantages (e.g., ability to identify and access managers with high alpha potential);
- views on strategic investment themes (e.g., overweight emerging markets within global equities);
- relative valuation differentials;
- practical constraints (e.g., availability of open or fund-raising high-quality managers in areas of interest); and
- the desire to further diversify exposures.

Within the framework we seek to measure and monitor implementation risks in a way that improves our ability to understand the risk/return trade-offs of implementation decisions at each level of decision making. Our analytics address how much active risk is being taken, how correlated these risks are, how these decisions affect equity beta exposure, and how the exposure affects total portfolio volatility. Given the changing nature of implementation risks and potential rewards, the framework uses a dynamic approach to understand risk characteristics of portfolio positions in the current environment and put them in a historical context.

INVESTMENT POLICY

A top-down statement that clearly expresses your most timeless and fundamental objectives and constraints. It reflects desired role-in-portfolio exposures (e.g., diversified growth, deflation hedge, and inflation sensitive), value-added performance objectives, and common risk factors of equity beta, illiquidity tolerance, and foreign currency risk. The policy is designed to provide those implementing portfolios with the appropriate guidelines for meeting long-term objectives.

LONG-TERM INVESTMENT PORTFOLIO (LTIP)

This refers generally to the investment portfolio, pool, or fund that is long term in nature and is being evaluated within the context of the Risk Allocation Framework. This may include, but is not limited to, a long-term endowment or foundation portfolio, private family investment pool, pension fund, insurance fund, liability settlement trust, or sovereign wealth fund. This term is used throughout this framework when referring to any long-term oriented investment portfolio under examination.

ROLE-IN-PORTFOLIO CATEGORIES

Investment policy within the Risk Allocation Framework classifies investments by the role they play in the portfolio. For endowments, foundations, families, and like investors that tend to have relatively high (e.g., 4%+) annual outflows from their LTIPs, this would include three broad categories: deflation hedge, inflation sensitive, and diversified growth (further detail provided below). Furthermore, these macroeconomic hedge allocations (deflation hedge and inflation sensitive) include only what we expect to be the purest means of protection.

Note that for investors, such as pensions or insurance companies, policy role categories would naturally differ. For example, at its simplest level, the policy portfolio for a defined benefit pension plan would involve allocating assets between a "hedging portfolio" that immunizes the pension liability and a "growth portfolio" that seeks to promote growth beyond funding the liability. The policy allocation categories may differ, but the philosophy behind them is the same. See Appendix A for a discussion of defined benefit pension plans.

Deflation Hedge. Deflation hedge assets are reflected at the policy level as high-quality sovereign bonds, which provide the most liquid and historically reliable protection for portfolios during periods of deflation/prolonged economic contraction and generally have benefitted from flight-to-quality sentiment. These assets are intended to support spending and other liquidity needs should the equity-like portion of the portfolio get hit in such environments.

Inflation Sensitive. There is a continuum of inflation-sensitive assets ranging from those that are more likely to provide inflation protection (provided they are reasonably valued), but offer a lower expected return and less opportunity to add alpha, and those that are less likely to protect against inflation, but have higher expected returns and provide more opportunity to add alpha.

Inflation-sensitive assets at the policy level include the former (commodities and short duration inflation-linked bonds) since we are seeking those investments that will best support liquidity needs and provide protection in periods of unexpected increases in inflation.

Diversified Growth. The diversified growth role category includes a mix of growth-oriented and diversifying (lower equity beta, but still relatively growth oriented) investments, which is represented, from a benchmark perspective, as global equities with a beta adjustment (see Equities with a Beta Adjustment above). In the policy setting process this allocation is set after policy allocations to deflation hedge and inflation-sensitive investments are set.

SHORT-TERM AND LONG-TERM RISK

While the precise definition of key short-term and long-term risks varies by investor, when evaluating risk tolerance, we look at the trade-off between long-term and short-term risk preferences. We define such risks generically, as described below, so the analysis is useful to investors universally.

Short-Term Risk. Risk of significant short-term declines in market value, expressed as the probability that the portfolio will decline by a meaningful percentage (e.g., 30%) during a short (e.g., five-year) period under stressed conditions. We use our stressed performance assumptions in this analysis, which reflect more volatile conditions than do our equilibrium return assumptions.

Long-Term Risk. Risk of failing to preserve purchasing power over the long term, expressed as the probability that the portfolio will fail to earn what it spends in real terms over a long time horizon (e.g., 20 to 25 years).

VALUE ADDED

Value added is a general term that refers to the outperformance of a portfolio relative to its benchmark. It is the simple arithmetic difference of returns. You can calculate value added at many levels ranging from manager returns less manager benchmarks to total portfolio returns less the policy benchmark. Some examples of value-added levels are defined below.

Implementation Value Added. Performance attributable to implementation (e.g., manager selection, manager structure, asset allocation detail) relative to policy portfolio benchmark returns.

Manager Selection Value Added. Performance attributable to your ability to select active investment managers that outperform their respective passive benchmarks.

Manager Structure Value Added. Performance attributable to the asset class substructures (capitalization, geography, style biases, etc.) relative to policy asset class returns. For example, your U.S. equity manager structure might reflect a higher small-cap allocation than exists in the U.S. equity benchmark. Similarly, a portfolio overweight to emerging markets relative to that of the global equity benchmark might be considered manager structure value added.

VOLATILITY-EQUIVALENT SIMPLE STOCK/BOND PORTFOLIO

A portfolio comprising only stocks and bonds, used as a simple way of expressing market returns at a given level of volatility with no value added through implementation. ■

Arlington Beijing Boston Dallas London Menlo Park Singapore Sydney