MEKETA

The Total Portfolio Approach for US Institutional Investors

CONTRIBUTORS
Frank Benham, CFA, CAIA
Colin Bebee, CFA

WHITEPAPER | APRIL 2026

The strategic asset allocation ("SAA") model, with its fixed targets across discrete asset classes, has dominated institutional decision-making for decades. Yet some of the world's largest long-term investors, including Canada Pension Plan Investment Board, New Zealand Superannuation Fund, Singapore's GIC, and Australia's Future Fund, have adopted an alternative framework called the Total Portfolio Approach ("TPA"). In 2025, CalPERS, the largest US public pension fund, began evaluating TPA for investing its investment portfolio. This has led many US-based institutional investors to reckon with a fundamental question: whether the governance, operational, and analytical features of TPA warrant a shift in how they think about governance and portfolio construction.

This paper is written for US-based institutional investors who have heard of TPA and want to understand it more fully. It helps define TPA, where it came from, how it differs from existing decision-making frameworks, and what its adoption would require in practice. The analysis focuses primarily on large public pension plans, as they represent the largest potential adopter base. However, TPA can be applied to virtually any type of portfolio with the necessary combination of internal and external resources. TPA has genuine appeal, but also material complications, and reasonable fiduciaries may reach different conclusions about fit.

TPA is often described as an alternative to SAA, offering increased flexibility and greater transparency. In practice, the key change is governance: what investment authority the board is delegating, how that authority is managed and monitored, and how the investment decision-making framework changes as a result.

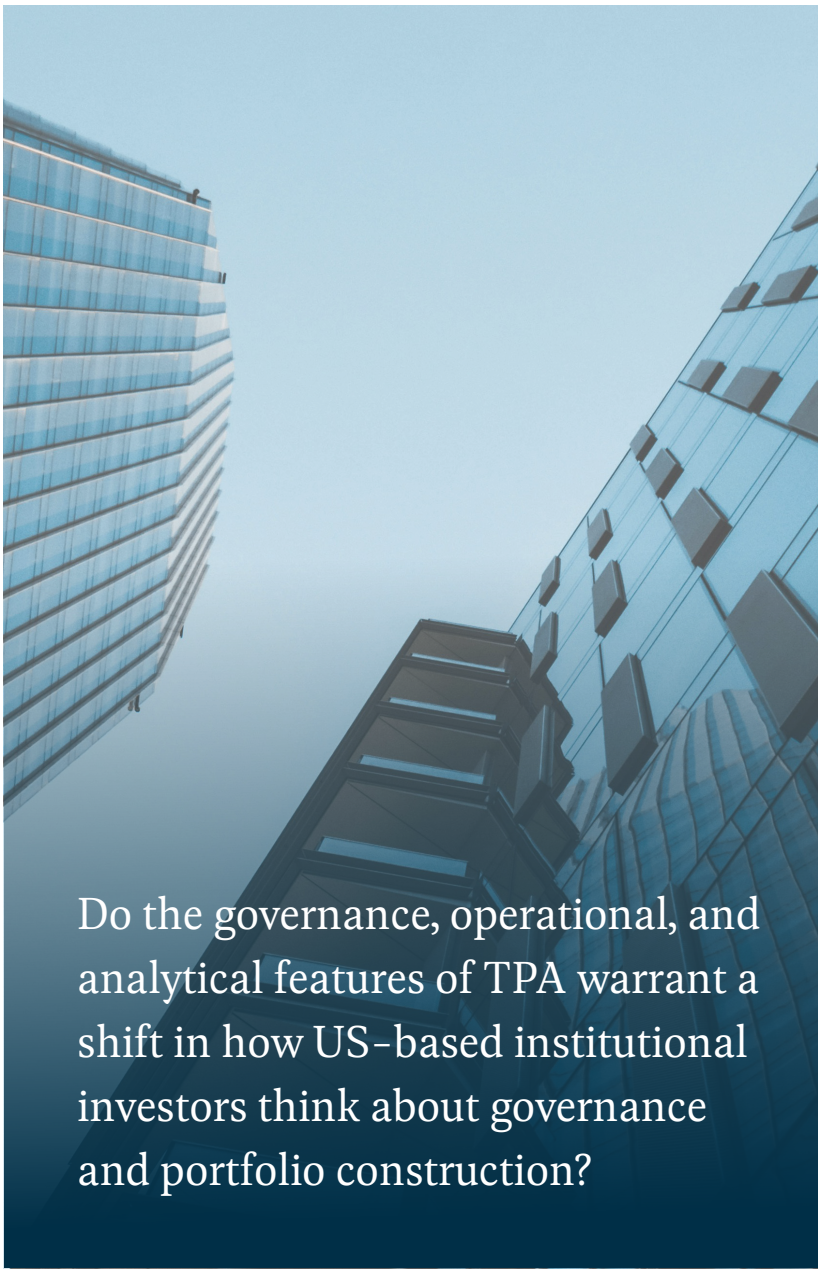

Key Takeaways

- **TPA offers positive attributes alongside material challenges.** While TPA provides compelling advantages for long-term portfolio management, it introduces meaningful complexities and demands careful consideration of new risks and operational requirements. TPA requires sustained organizational commitment to implement and manage effectively.
- **First and foremost, TPA is a governance model.** Its most consequential feature is the transfer of capital allocation authority from a governing body to an investment team. Under SAA, the board approves a strategic asset allocation with fixed targets and ranges and relies on those targets to control risk. Under TPA, the board typically approves a simplified reference portfolio and an active risk budget, transferring the granular capital allocation decisions to the investment team.
- **TPA is best understood as a spectrum.** In its mildest form, TPA judges every investment by one yardstick: what does it do for the whole fund? In its bolder forms, it rewires decision-making so that boards set portfolio-level objectives and risk tolerances while professional investment teams dynamically allocate capital across asset classes. Major practitioners treat TPA not as a single, off-the-shelf model, but as a governance philosophy that institutions customize to their own operational capacity and fiduciary structure.
- **The reference portfolio replaces the SAA as the primary performance benchmark.** In practice, the reference portfolio is a simple, liquid, publicly investable portfolio, often composed of a weighted combination of global equities and investment grade bonds (e.g., Treasuries), that defines the investor's risk appetite. Every investment decision is judged by whether it improves the expected risk-adjusted returns that the reference portfolio could deliver.
- **The operational requirements are substantial as are the changes in organizational culture.** TPA demands integrated risk measurement across public and private markets, a common factor language for the entire portfolio, and data infrastructure that most institutions either do not have or would need to materially enhance. Moreover, TPA requires a cognitive and cultural shift among investment teams to evaluate new and existing investments *across* asset classes rather than simply *within* asset classes.
- **Performance claims for TPA should be treated with caution.** Some industry studies have found that TPA practitioners outperformed SAA users. However, these studies carry significant methodological limitations, including small samples, large-institution bias, and benchmark-choice problems that make direct attribution to TPA governance difficult. Investors should not adopt TPA on the basis of historical performance alone.
- **Many of TPA's best practices can be adopted without full conversion.** Investors who are not ready, or not suited, for the complete model can still benefit from evaluating risk factors alongside asset class exposures, applying total-fund active risk limits (i.e., risk budgets), building cross-asset class team collaboration, and treating every allocation as competing for a share of a common risk pool.

The Problem TPA Is Trying to Solve

The Architecture of Strategic Asset Allocation

Since Harry Markowitz formalized mean-variance optimization in 1952, strategic asset allocation has dominated institutional portfolio construction.[1] The governing body (typically a board of trustees), often advised by consultants and actuaries, selects target allocations across discrete asset classes, such as public equities, fixed income, real estate, and private equity, based in part on forward-looking long-term return and risk assumptions. The portfolio is then divided into silos, and investment teams manage each asset class against the respective benchmark approved by the board. The board tends to revisit the allocation no more often than annually and no less often than every five years.

Boards typically delegate day-to-day investment decisions to professional investment teams who operate within the board's strategic framework. These investment teams may be external, internal, or both, with larger investors tending to opt for greater use of internal staff. In theory, this division is clean: the board sets policy and the investment team executes within bounds. In practice, the boundary between strategic and tactical decision-making may be permeable, creating structural tensions.

Where the Cracks Appeared

While SAA has served many institutions well, a number of systemic issues have become apparent. It is worth noting that some critiques conflate challenges with the SAA framework itself with problems specific to mean-variance optimization, which many SAA practitioners have long moved beyond.

- **Silo Incentives.** People respond to incentives. If investment decisions are evaluated against an asset class benchmark, that benchmark will dominate decision-making, sometimes at the expense of the total portfolio. A fixed income allocation, for example, might be intended to hedge equity downside risk, but a team measured against a fixed income benchmark may tilt toward credit risk to improve their relative return, inadvertently increasing, rather than reducing, the portfolio's equity sensitivity. Silo incentives can also produce over-diversification or "benchmark-hugging" behavior within asset classes, as teams spread investments broadly to minimize benchmark-relative risk rather than concentrate in their highest-conviction ideas.

- **Hidden Risk Exposures.** A siloed approach can create a false sense of diversification, where each asset class portfolio appears diversified in its own right, potentially masking the fact that assets in multiple asset classes are exposed to the same underlying risk drivers. For example, a private equity portfolio heavy with turnaround buyout funds, a hedge fund sleeve full of distressed credit strategies, and a private debt book with debt-for-control positions may sit in three separate SAA buckets, each with its own benchmark, yet share essentially the same underlying exposure to distressed credit conditions. Under a siloed framework, these overlapping exposures may go unrecognized until a downturn reveals it.

- **The Rebalancing Problem.** SAA assumes the portfolio can be rebalanced to target weights or, more often, to within target ranges, as markets move. However, in practice, private markets cannot be rebalanced on demand. When public equity valuations fell sharply in 2022 and many private market valuations were not marked down, many funds found themselves overweight private markets at precisely the moment they might have wanted to increase public equity exposure. Given that public and private equities are often treated as the closest proxy for each other (from risk exposure and economic driver standpoints), decision-makers might have felt they could not increase public equities while they were above their target range for private equities. Hence, target allocation bands around illiquid assets can constrain flexibility in rebalancing.

- **Benchmark Gaming.** When investment teams are rewarded for outperforming benchmarks, there is an incentive, not always conscious, to select asset class benchmarks that are easier to beat. Board members are often not intimately familiar with the detailed composition and risk exposures embedded within each underlying benchmark. Hence, the number of asset classes and proliferation of benchmarks under SAA amplifies the opportunity for this dynamic.

What TPA Actually Is

The Core Idea

TPA's central premise is that the portfolio should be managed as a single unit rather than as a collection of separate asset classes. Every investment decision, whether to commit to a new buyout fund, tilt toward value equities, or increase interest rate exposure, should be evaluated by its marginal contribution to the whole portfolio's risk and return, not by its expected performance relative to an asset class benchmark. While this concept is also present under an SAA framework, the TPA explicitly places this goal front and center.

In the words of CalPERS CIO Stephen Gilmore, who came to CalPERS after running the New Zealand Superannuation Fund (NZ Super): "At its most basic, a total portfolio approach asks the team to construct a portfolio to achieve the objective. It is direct. So, everything is focused on the whole of the portfolio."[2]

The Reference Portfolio

The most visible structural feature of TPA is the reference portfolio. Where SAA produces a policy portfolio with fixed targets for each asset class, TPA replaces that with a benchmark that is intended to represent the fund's risk appetite and return expectation using an easily replicable portfolio composed of a mix of two liquid asset classes, either stocks and bonds or (less commonly) stocks and cash.

The reference portfolio is not intended to represent the actual portfolio. Rather, it serves as a baseline that reflects the governing body's risk level. The reference portfolio shows the returns that may be achieved at the respective risk level, without employing any active strategies or complex investment decisions. Every active decision (defined as any allocation outside of the reference portfolio) made is then evaluated against this baseline. If a commitment is made to a buyout fund, a TPA investor is implicitly deciding that the buyout fund will outperform what the reference portfolio would have returned over the same period on a risk-adjusted basis, net of all costs. Note that a simpler benchmark (i.e., the reference portfolio) does not guarantee that a simpler portfolio is employed. Complexity often migrates from benchmark design (under board control) to implementation (e.g., overlays, leverage, liquidity management, and private market pacing).

At NZ Super, for example, the reference portfolio is an 80/20 equity-bond split, fully hedged back to the New Zealand dollar, and passively managed.[3] At CalPERS, the newly adopted reference portfolio is a 75/25 global equity-bond mix, replacing a policy benchmark previously composed of eleven separate asset class benchmarks.[4] Reference portfolios may change over time as governing bodies turn over and as capital markets change, but the fundamental premise remains consistent.

The reference portfolio serves two purposes simultaneously. First, it defines risk appetite and return expectations: a more equity-heavy reference portfolio implies a higher tolerance for risk and a higher expected return. Second, it creates a clear measure of whether active management is adding or detracting value at the total fund level. This assessment is harder to obscure than asset class-specific benchmarking. For example, the relevance of portfolio versus benchmark comparisons over shorter time periods can be significantly impaired.

Done well, the reference portfolio can clarify opportunity cost for every investment. One potential advantage of setting a reference portfolio is that it can establish a stable risk appetite. This approach helps guard against pro-cyclical behavior that could arise from changes in board membership or shifts in return targets as capital market assumptions evolve. Done poorly, the fund may appear to generate "added value" simply by holding its normal allocation and benefiting from taking on more risk, or simply different risks, than the benchmark. The choice of a reference portfolio is a consequential one and deserves at least as much scrutiny as the selection of any other benchmark. It involves several difficult judgments: whether to calibrate the reference portfolio to the fund's historical risk level or its forward-looking expected return target (or a point in between the two); how to adjust when capital market assumptions change; and whether observed volatility in a portfolio containing smoothed private market valuations reliably guides the appropriate equity weight in a public-markets-only reference portfolio.

Moreover, modern financial theory indicates that there are multiple systematic risk factors that reward investors for simple participation. Reference portfolios in practice are largely focused on equity and interest rate risk, with the latter potentially excluded. This implies that outperformance compared to a reference portfolio may be achieved by simply participating in other markets (e.g., credit) that would be captured in an SAA framework. There are no consensus answers for how best to structure a reference portfolio, and different TPA adopters have made different choices.

The Active Risk Budget

Alongside the reference portfolio, TPA funds typically adopt an active risk budget, expressed in basis points of tracking error (i.e., active risk) at the total fund level.[5] This risk budget quantifies how far the actual portfolio can deviate from the reference portfolio in terms of expected volatility of excess returns. CalPERS, for example, has adopted an active risk limit of 400 basis points, with an expected operating range of 250 to 350 basis points.[6] Many practitioners pair a hard limit with an internal operating band and pre-agreed triggers. Approaching the limit may prompt a discussion about which risks to reduce, rather than reopening a policy debate in the middle of a market episode.

Risk budgeting is not uncommon among large US institutional investors, but it is often applied at the asset class level or focused on liquid markets composites only, rather than at the total portfolio level. Under TPA, every new investment is assessed in part by its incremental contribution to the risk budget. An investment that would comprise a large portion of the risk budget has to justify that allocation relative to the available alternatives, including existing positions. Proponents of TPA describe this as forcing investment ideas to compete for capital.

Singapore GIC's implementation of TPA offers a useful illustration of how this competition for capital may work in practice. GIC frames the portfolio as built from two parts: a policy portfolio (representing beta, or systematic risk exposures) and an active portfolio (an overlay of alpha exposures). Every alpha strategy is funded from the sale of a slice of the policy portfolio, meaning each strategy must deliver risk-adjusted excess return over its "funding cost" (i.e., the risk-adjusted return it would have earned had the capital remained in the passive policy portfolio).[7] This specific architecture makes the opportunity cost of every active decision explicit and measurable.

Factor Language

One of TPA's more technical features is its use of risk factors as a common language for measuring and comparing risk across asset classes. Rather than (or in addition to) reporting that private equity is 17% of the portfolio and real estate is 10%, a TPA fund might report that its portfolio has a 70% equity beta equivalent, 15% rates exposure, and 10% credit spread exposure, cutting across all asset classes simultaneously. (The mechanics of how factor exposures are estimated, including for private market assets, are discussed in the Implementation Mechanics appendix).

This cross-asset factor decomposition can help identify hidden risk exposures at the level of the factors chosen. It is worth noting, however, that standard factor frameworks are not designed to capture more granular or thematic exposures. For example, a portfolio heavily exposed to AI-related capital expenditure or the risk of AI-driven industry disruption would not necessarily reveal that exposure through a conventional factor decomposition. Factor-based risk measurement is a useful lens, but not a complete one. As such, it is effectively implied that other forms of risk exposure are evaluated by the investment teams, no different than how siloed asset class teams often examine similar exposures.

The need to measure factor exposure is partly what makes TPA operationally demanding. Computing factor exposures across a portfolio that includes both public and private market assets requires sophisticated risk systems and reliable underlying data. The challenge extends to strategies (e.g., hedge funds) that may not provide full position-level transparency, requiring the use of proxy-based estimates that may not accurately reflect true exposures.

Different TPA adopters have made different choices about which factors to include; most are deliberately parsimonious. Canadian Pension Plan (CPP Investments), for example, uses growth, rates, and spreads. Some non-US investors include currency as a meaningful factor given the size of their currency tilts. The choice of factors is both a technical and a judgment call, and different choices will produce different pictures of the portfolio.

Figure 1
Key TPA Adopters and Their Structures

Fund	AUM (USD Approx)	Reference Portfolio	Year TPA Adopted
NZ Super (New Zealand)	~$50 bn	80/20 equity/bonds	~2009
CPP (Canada)	~$570 bn	85/15 equity/bonds[8]	~2006
Future Fund (Australia)	~$270 bn	CPI plus 4% to 5%[9]	~2010
GIC (Singapore)	~$800-900 bn	65/35 equity/bonds[10]	~2022
CalPERS	~$556 bn	75/25 equity/bonds	2026[11]

Sources: NZ Super: fund annual report and nzsuperfund.nz; CPP Investments: CPP Investments annual report and investment statement, April 2025; Future Fund: Future Fund annual report; GIC: GIC ThinkSpace, August 2022; CalPERS: CalPERS board materials, November 2025.

A Few Case Studies

The performance claims most commonly cited for TPA derive from a handful of institutions that have practiced it over long periods. Examining these cases carefully is more useful than citing aggregate statistics.

New Zealand Superannuation Fund

NZ Super is the institution most often cited as the proof of concept for TPA. It is unusually transparent, has practiced TPA longer than most, and has a clear reference portfolio against which its results can be measured. The fund's reference portfolio is an 80/20 passive equity-bond benchmark, with all foreign currency exposure hedged back to the New Zealand dollar. Over the long term, NZ Super has generated returns above its reference portfolio, though the magnitude understandably varies by measurement period.

Institutions should be cautious about over-interpreting this performance. There are unique aspects to NZ Super's performance that may not be easily replicated by others. Much of NZ Super's historical active value-add has been generated by strategic tilting, or dynamic asset allocation across more than 30 markets. Their approach is designed to be systematic, relying on mean reversion and volatility harvesting, with a focus on assets/decisions with a high information ratio. This includes active currency management, which in NZ Super's case, includes decisions about how much of the fund's large foreign currency exposure to leave unhedged.[12] Arguably, these activities rely heavily on judgment, timing, skill, and perhaps luck. Therefore, investors should avoid assuming that copying NZ Super's TPA framework will yield similar outcomes. Investors should evaluate whether their own organizational capabilities and risk tolerances would be conducive to similar outperformance.

Canadian Pension Plan

CPP, Canada's $570 billion national pension plan, is among the earliest and most sophisticated practitioners of TPA. Its approach emphasizes a total fund perspective on risk and a competition-for-capital framework in which each new investment opportunity must justify its inclusion by improving the portfolio's overall risk-return profile.[13]

CPP has historically used leverage, a degree of freedom that TPA permits but SAA frameworks tend to constrain through policy bands. The fund has produced strong absolute returns, though its performance relative to its reference portfolio has been less impressive.[14]

CPP's investment framework also illustrates how factor language can be institutionalized formally. The fund views factors such as economic growth, rates, and credit spreads as compensated over longer horizons and assesses all portfolio investments through this lens. Its formal investment statement describes board-approved market risk limits while granting management discretion to make and implement investment decisions within those limits.[15]

The Mechanics of Implementation

Governance: Who Decides What

TPA's most fundamental change is governance, not analytics. It moves capital allocation authority from the board to the investment team. Under TPA, the board sets the risk appetite (through the reference portfolio and the active risk budget). The investment team decides, within that budget, where to allocate capital.[16]

This delegation has genuine appeal. In the US, boards are not typically composed of investment professionals, meet infrequently, and contend with a wide variety of non-investment issues. Board turnover can disrupt institutional memory, shift priorities, and introduce abrupt changes in risk appetite. Hence, it is plausible that a dedicated investment team may simply be better positioned to make ongoing allocation decisions.

Risk Measurement

Risk measurement is the technical heart of TPA, and it is considerably more challenging than its proponents sometimes acknowledge. Managing assets within a TPA framework requires knowing, in near-real time, how much risk the portfolio is taking. This includes both total fund risk (the overall volatility and drawdown potential of the portfolio) and active risk (the degree to which the portfolio deviates from the reference portfolio). TPA proponents often, justifiably, fall back on their long-term investment horizons and how performance/risk comparisons should align with such timeframes.

For public markets, commercial risk systems can compute factor exposures, tracking error, and marginal risk contributions with reasonable accuracy. For private markets, the challenge is considerably harder and has two dimensions. The first is staleness: private equity, real estate, infrastructure, natural resources, and private credit positions do not mark to market daily, and their reported valuations lag by a quarter or more. When a risk budgeting tool computes total fund tracking error using stale private market valuations, the output may substantially misestimate risk. The second is the choice of benchmark: private markets will show very high observed tracking error relative to a liquid reference portfolio. The standard solution is to map private market assets to a public market equivalent and compute active risk relative to that proxy.

The data problem is equally acute. Visibility into the factor exposures of individual private market positions requires data that general partners do not routinely provide in standardized form. Aaron Filbeck of CAIA describes the situation plainly: "You can throw technology at a problem, but if you don't have underlying information and data, that makes it really challenging."[17]

Tactical Flexibility and Its Limits

One of TPA's most appealing features is the freedom it provides to make tactical shifts or opportunistic investment. Under a SAA with narrow policy bands, an investment program is constrained in how much risk can be added or removed from a portfolio.

TPA removes this constraint, at least within the overall risk budget. Investment teams can shift the portfolio dynamically and deploy capital into opportunities as they arise. However, decision speed and flexibility are a double-edged asset. A shorter governance process can help a fund buy into dislocations, but that presumes the investment team is correct in identifying and timing them. Likewise, greater flexibility can amplify both gains and losses.

In practice, the active risk constraints under SAA can often be behavioral rather than formal. That is, the policy bands may be sufficiently wide to provide adequate flexibility, but the investment team tends not to fully use this latitude. Adoption of a formal risk budget (with an active risk target, not just ranges) may provide greater clarity on the expectations for the use of active risk.

The question is whether this flexibility produces better outcomes in practice. Large market dislocations, like that which occurred during the Global Financial Crisis, occur infrequently, making it hard to identify whether an investor has sustainable skill in identifying and timing such dislocations. Tactical asset allocation has, at best, a mixed track record. Academic and practitioner literature consistently finds that macro market timing (i.e., moving between asset classes based on near-term relative performance views) is difficult to execute profitably over time, with a wide dispersion of outcomes across managers and cycles.[18] The freedom to be dynamic is only as valuable as the judgment guiding its use, and that judgment is difficult to assess in advance.

The Objectives and Concerns

TPA is Not The Only Solution

The most substantive critique of TPA is that many of the problems it addresses can be tackled within a modernized SAA framework, without the governance disruption that a full TPA transition entails. A sophisticated fund can impose total-fund risk budgets without abandoning asset class benchmarks. It can run cross-team investment committees to catch hidden risk exposures. It can use risk factors as an overlay to identify and monitor exposures. It can create an opportunistic bucket for cross-asset ideas that do not fit neatly into existing silos.

Hence, investors might consider taking an a la carte approach to TPA, selecting the portions that they like. As Texas Teachers Retirement System's (TRS) investment management division concluded in early 2026: "Traditional SAA and TPA are different mixes of the same ingredients which result in different prioritization."[19] For some institutions, a hybrid approach may provide the best solution.

Centralization and TAA

TRS CIO Jase Auby has articulated a more philosophical objection: that TPA represents a top-down, centralized decision-making process that is potentially "brittle" because it relies on one person or one small team. He contrasts this with TRS's approach of pushing authority deeper into the organization, with autonomy within each area well-defined and supported by good risk controls.[20] A TPA practitioner might argue that TPA should be "joined up" from a top down and bottom up perspective, with ideas from the asset class teams affecting top down decisions.

Auby also raised a pointed challenge to tactical asset allocation (TAA), which is often treated as a centralized, whole-portfolio tool within TPA implementations. He noted that TRS terminated its own TAA activities in 2017 after years of unsatisfactory results, characterizing TAA as "low breadth" decision-making: few, highly consequential choices, compared with the many decisions involved in security selection.[21] The intuition is straightforward and is the core of "The Fundamental Law of Active Management":[22] with a large number of independent decisions, a manager only needs to be right slightly more often than wrong for skill to compound into meaningful outperformance. Where security selection benefits from large numbers of decisions and the law of large numbers, TAA bets tend to be few and, because they are often driven by the same macro environment, may be more correlated with one another than the breadth count alone would suggest.

This critique deserves serious weight. The pursuit of portfolio coherence can come at the cost of specialization. The institutional investment industry has developed deep asset class expertise for good reason: specialists with domain knowledge may consistently outperform generalists making top-down calls, particularly in less efficient markets. TPA discussions sometimes underplay this trade-off.

The Evidence is Thin and Open To Misinterpretation

The empirical case for TPA's superior performance is built on small samples, large-institution bias, and methodological choices that make direct attribution to TPA governance difficult (see the appendix). There is also a concern that TPA funds' outperformance reflects, at least in part, a systematic redefinition of their benchmarks to something easier to beat.

Risk Measurement Remains a Challenge

The claim that TPA provides a clearer, more integrated view of portfolio risk rests on an assumption that risk can be measured adequately across both public and private market assets. Assumptions around how to treat private markets have long challenged investors who engage in risk budgeting at the total fund level. For example, risk budgeting with private markets typically forces investors to choose between two viewpoints: economic reality vs. valuation-based pricing. Moreover, the factor frameworks that TPA practitioners use, while conceptually sound, involve assumptions about how private market assets map to public market risk factors that are difficult to validate empirically, particularly during periods of market stress when those assumptions are most likely to break down.

The Cost of Conversion

One question that performance comparisons and governance discussions tend to sidestep is the cost of conversion itself. Building TPA-ready infrastructure (e.g., integrated risk systems capable of spanning public and private markets, a factor framework with reliable data feeds, redesigned performance measurement structures, and the governance documentation that supports meaningful board delegation) is neither cheap nor fast. For large funds with substantial internal teams, some of this infrastructure may already exist in partial form, making the marginal cost of conversion more manageable. For most US institutional investors, however, the build-out would represent a meaningful multi-year investment in systems, data, and organizational capability, with costs that would need to be recovered through incremental alpha before the governance change generates net value. This is not an argument against TPA, but an argument for treating adoption as a capital allocation decision in its own right, subject to the same cost-benefit discipline that TPA itself demands of every other investment. Additionally, large investment organizations are not without internal frictions, and the required change in culture is an important consideration.

The Impact of Turnover

Leadership transitions pose a particular risk. For example, a change in CIO, particularly if the outgoing leader was the primary advocate for TPA, can disrupt established relationships and jeopardize the continuity essential to the model's success. Similarly, board turnover introduces uncertainty, as new members may challenge the rationale for TPA, especially during stretches of poor portfolio performance.

Governance Alignment is Essential

A critical, often underappreciated barrier to TPA adoption is whether the board and the investment team share a sufficiently clear and common understanding of objectives, risk tolerance, and decision-making authority. If the board lacks confidence in the team's judgment, expertise, or alignment with the fund's objectives, it is unlikely to endorse a governance model that shifts significant discretion and accountability to that team. Building and sustaining that confidence requires clear communication, demonstrated competence, robust risk controls, and transparent reporting, not as a one-time threshold to clear, but as an ongoing condition of the delegation relationship. Boards need assurance that the investment team can manage TPA's complexities, especially in challenging markets or after periods of underperformance.

Governance and Accountability are Complicated

Fundamentally, there are two key investment decisions: asset allocation (a Board-level decision under SAA) and implementation (a staff-level decision for sophisticated plans with delegated authority, which would represent the primary candidates for exploring TPA). TPA changes how accountability works: while SAA allows for clear performance attribution at the team or asset class level (e.g., the real estate team either beats its benchmark or it does not), TPA consolidates accountability at the total fund level.

This shift can theoretically align incentives and discourage gaming of individual benchmarks, but it also makes it more difficult to attribute good or bad outcomes to specific decisions, as total fund results are heavily influenced by factors such as market beta, the equity tilt of the reference portfolio, or other external conditions outside the team's control. As a result, boards that adopt TPA without addressing the challenges of monitoring, reporting, and communication invite scrutiny that can be difficult to address, not because the decisions were wrong, but because the framework itself makes them harder to attribute.[23] TPA can be straightforward to explain conceptually but difficult to defend after a bad year, when stakeholders (e.g., legislators, participants, donors) demand a clear narrative about why specific choices were made and how controls operated.

Who Should Consider TPA?

TPA is not appropriate for every institution. Several prerequisites appear to be necessary conditions rather than nice-to-haves. Note that external solutions (e.g., OCIOs, specialist consultants) may be able to substitute for internal capabilities in some of these areas.

Scale. The benefits of TPA are most accessible to large investors. The analytical and operational infrastructure required to implement TPA properly has substantial fixed costs. Smaller funds may find those costs difficult to justify, and their investment teams may lack the breadth of expertise needed to evaluate cross-asset opportunities. However, smaller teams may find it easier to collaborate.

Investment Management Capability. TPA works best when the entity managing the portfolio is capable of making allocation decisions across multiple asset classes without operating within siloed mandates. Organizations that rely primarily on external managers operating independently within their own mandates may find it difficult to implement the competition-for-capital logic at the heart of TPA.

Risk and Reporting Infrastructure. Successful TPA implementation requires the capability to construct and oversee a total fund risk budgeting system, measure and monitor factor exposures across public and private market assets, and generate reporting that gives boards and stakeholders meaningful visibility into how risk is distributed and managed. This infrastructure must be built before it is needed, not assembled in response to a market event.

Governance Sophistication. TPA requires a board willing to set a risk appetite and genuinely delegate, and an investment team capable of accepting portfolio-level accountability. Both are harder than they sound. Boards naturally want to understand what they own, and asset class reporting provides a familiar vocabulary. Risk-budget reporting, equity beta equivalent, liquidity forecasts, and many other analytical components of TPA, require a different kind of financial fluency than many boards may be accustomed to.

Governance Clarity. For US institutional investors, governance requires clarity about lines of accountability beyond the board itself (e.g., to plan participants or elected officials). A TPA model that works well in a sovereign wealth fund context, where governance lines are relatively clean, may face additional friction where multiple stakeholders have legitimate claims on the fund's decision-making.

Long Time Horizon. TPA's competition-for-capital model often favors illiquid, long-duration assets, which is appropriate for a perpetual fund but may create complications for a fund facing substantial near-term liquidity needs. Such investors (e.g., mature pension plans with many retirees relative to active members) need to be thoughtful about how much risk they can take with their asset mix, though this is not unique to TPA. The horizon for reviewing success/shortcoming also increases. For example, if active decisions result in higher private markets allocations, even a period as long as five years may not provide much insight on the comparative results.

The SAA/TPA Continuum

For institutions that find TPA appealing, the most practical near-term approach may be a selective incorporation of TPA principles within an existing SAA framework. Indeed, investors may want to think of a continuum between traditional SAA and full TPA, and choose the point on that spectrum that suits them best. Such a hybrid approach could include:

- **Adding a total-fund risk budget** alongside (though not necessarily replacing) asset class policy bands. Policy ranges could also be widened to create more flexibility. Texas Teachers has done this, demonstrating that the two approaches are not mutually exclusive.

- **Requiring total-fund risk attribution for significant investment decisions.** Any investment proposal above a defined size threshold would include an analysis showing how the proposed investment affects total fund tracking error, factor exposures, and liquidity profile. This analysis could be produced by the investment team or a dedicated risk function.

- **Designating an opportunistic allocation** with a range of 0% to 5-10% of the portfolio. By setting a range rather than a target, the allocation can be genuinely opportunistic. It would ideally be measured against the total fund policy benchmark (or a reference portfolio) rather than a single asset class benchmark.

- **Adopting a common risk factor framework** across asset class teams. This means agreeing on a shared set of risk factors and requiring all teams to report their exposures in those terms. A common language makes aggregation of risk exposures possible.

- **Utilization of a reference benchmark** in combination with a policy benchmark at the total fund level. The reference benchmark serves as a gauge of asset allocation success (reflecting input from both the governing body and internal/external resources), whereas the policy benchmark focuses more on implementation success based on the board-approved asset allocation.

Institutional investors can introduce these steps gradually, avoiding the governance disruption associated with a full TPA conversion and without requiring the operational infrastructure demanded by a comprehensive TPA model. Since many of these steps are part of the pathway toward TPA, they can also serve as transitional measures, allowing organizations to assess their readiness before committing to full implementation.

Conclusion

TPA represents a meaningful evolution in how institutional investors can think about portfolio construction and governance. Its core insight, that every investment should be judged by what it contributes to the whole portfolio rather than how it performs against an asset class benchmark, is sound, and the structural problems it addresses are real. This focus on total fund results is fully aligned with the overarching objectives of institutional investors (e.g., retirement benefit payments, endowment spending, etc.). Siloed decision-making, hidden risk exposures, benchmark gaming, and the growing complexity of modern portfolios are genuine challenges, and TPA offers a coherent framework for addressing them. But coherence is not the same as simplicity, and adoption is not the same as success.

What TPA changes, most fundamentally, is not the portfolio; it changes who makes decisions, on what basis, and how accountability is structured. That governance shift has genuine appeal, particularly for institutions whose boards struggle to make timely, portfolio-level decisions within the constraints of a rigid asset class structure. It also carries real risks: concentrated decision-making authority, greater reliance on tactical asset allocation, and the challenge of communicating outcomes to stakeholders who are accustomed to the familiar vocabulary of asset class reporting. Institutions considering TPA should spend at least as much time on these governance questions as on the analytical ones. Remedies to many of these challenges exist, but their relevance and utility will depend on the given institution.

The operational demands are equally significant. Integrated risk measurement across public and private markets, reliable factor exposure data, and the technology infrastructure to support total-fund portfolio management represent a meaningful investment of time, money, and organizational focus. For most institutions, that build-out is a multi-year undertaking, and its costs need to be weighed honestly against the expected benefits before a governance decision is made.

For many US institutional investors, the most practical path may not be full TPA adoption but a deliberate incorporation of its principles into an existing framework. The competition-for-capital logic, a common risk factor language, a total-fund risk budget alongside asset class ranges, and an opportunistic allocation evaluated against a portfolio-level benchmark are all steps that can improve investment decision-making without requiring a complete governance overhaul. TPA is, in the end, less a destination than a direction, and most institutions can take something positive from evaluating its characteristics.

Appendix: Fit Assessment for US Institutional Investors

The following questions can help boards and investment teams assess fit without falling into an "all-or-nothing" framing. For many institutions, the honest answers may point toward something in between (e.g., a hybrid approach) rather than a full TPA conversion.

If the answers are mostly "yes," TPA, or a serious move toward it, is worth exploring in depth, with full attention to implementation sequencing and the governance changes required.

If the answers are mixed, a hybrid path is likely the right near-term answer. An investor can adopt the elements of TPA that address identifiable problems, without completely disrupting the existing governance structure.

If the answers are mostly "no," the existing SAA framework may be the more appropriate governing structure. That is not a failure; it is an honest assessment of where the institution is. Many of TPA's best practices, such as thinking in risk factors, identifying hidden risk exposures, and thinking about the impact on the overall portfolio, can be grafted onto an SAA framework without a full governance overhaul.

Question	What "Good" Looks Like
Can trustees agree on a small number of fund-level objectives and constraints?	Clear priorities among return, solvency, liquidity resilience, and mission, ranked and agreed upon
Can the board define and defend a reference portfolio or policy anchor?	Transparent, investable, and understood by all stakeholders including external ones
Does the organization have (or can it access) portfolio-level risk tooling?	Regular scenario analysis, liquidity forecasting, and clearly defined risk limits across the whole portfolio
Does the investment team have breadth and cross-asset integration capability?	A repeatable "competition for capital" process that works across the entire investment team
Is the governance culture comfortable with discretion and accountability?	Delegated authority with rigorous total-fund reporting and genuine board engagement on risk, not just returns
Can the organization communicate the model publicly and withstand scrutiny after a down year?	A stable "how we manage risk" narrative that does not depend on markets cooperating
Does the fund have sufficient scale to justify the infrastructure investment?	Meaningful internal investment management capability and AUM sufficient to spread those costs
Is the investment team stable enough to sustain the model through leadership transitions?	Documented processes, shared factor language, and institutional knowledge that survives any individual's departure

Appendix: The Performance Evidence

One study of large asset owners reported that TPA-governed funds outperformed SAA-governed peers by approximately 1.3% per year over a decade.[24] A comparison of leading TPA practitioners against endowment benchmarks have similarly shown meaningful outperformance in absolute terms, with the average 10-year return for four major TPA funds running roughly 140 basis points above the endowment study average over the same period.[25]

The evidence for TPA outperformance should be approached with considerable skepticism for two reasons. First is the methodological fragility of the studies. Both studies focus on a small number of funds, typically those that are among the largest and most sophisticated globally, and compare them against broader peer groups that include smaller, less well-resourced institutions. Factors such as scale, talent, governance quality, home-country investment advantages, preferred access to deal flow, internal management teams, co-investment capabilities, and negotiating power on fees likely explain a meaningful portion of the performance gap. TPA does not guarantee any of these potential advantages. Additionally, the TPA adopters in the studies tend to exhibit higher risk appetites, including larger allocations to growth assets, private equity, and illiquid strategies, as well as the use of leverage. These factors would be expected to boost returns in a period of sustained economic expansion, irrespective of governance structure. Without adequately controlling for these differences, attributing the performance gap directly to TPA governance is difficult to justify. Even CalPERS's CIO, when presenting such data, cautioned the board that the sample was "very small" and advised against relying too heavily on those findings.[26]

Beyond methodological concerns, comparisons against benchmarks warrant further skepticism due to potential gaming of the reference portfolio. TPA funds may appear to generate "value add" simply by adopting a reference portfolio that carries less exposure to risky assets (e.g., stocks, private equity, credit) than the actual portfolio. When those risk assets outperform, the difference between the actual portfolio and the reference portfolio can be mistaken for alpha, despite being largely a function of structurally holding more risk than the benchmark. This dynamic complicates the attribution of performance to skill or governance and underscores the need for caution when interpreting TPA-related outperformance claims within peer comparisons.

None of this means TPA underperforms SAA. Rather, it implies that the evidence for outperformance is not compelling enough to settle the question.

Appendix: Implementation Mechanics

This appendix is intended for investment teams who need to understand not just what TPA is conceptually, but how its core technical components might work in practice. It covers four building blocks: risk budgeting, factor exposure calculation, illiquidity adjustment, and currency management.

Risk Budgeting

The key feature of a TPA risk budget is that it is not pre-allocated by asset class. Instead, it is consumed incrementally as investment decisions are made. Each new active position, be it an overweight to private equity relative to the reference portfolio, a tilt toward value equities, a large infrastructure commitment, etc., contributes some amount of tracking error to the total. The investment team tracks this cumulative consumption and ensures the total stays within the budget.

When the portfolio approaches or breaches the active risk limit, the investment team is typically required to reduce active positions to bring the portfolio back within bounds. Most TPA funds pair a hard ceiling with an internal operating band and pre-agreed triggers. This might, for example, include a requirement to convene a risk review if tracking error exceeds the midpoint of the operating range and to reduce active positions if tracking error approaches the ceiling.

A simple illustration: suppose a fund has a 300 basis point tracking error budget. It makes three active decisions:

- A 5% overweight to global equities (relative to the reference portfolio) that contributes roughly 80 basis points of tracking error
- A private equity allocation that, after mapping to its public market equivalent, contributes roughly 120 basis points
- An infrastructure allocation that, again after mapping to its public market equivalent, contributes roughly 60 basis points

Those three decisions consume approximately 260 basis points of the active risk budget, leaving roughly 40 basis points for additional active positions. If the team wants to add a meaningful credit tilt, it either needs to reduce one of the existing positions or accept that it is approaching its limit.

In practice, the math is more complex because active positions are not independent, because their tracking error contributions interact based on correlations. Similarly, it is important to note that all ex-ante risk projections naturally incur measurement error, and all risk models are based, to some degree, on history. A private equity overweight and a public equity overweight, for example, share significant equity beta and therefore contribute less combined tracking error than the sum of their individual contributions. Risk systems handle this aggregation using covariance matrices, but the conceptual logic is the same: each decision consumes a share of a finite budget, and the team must prioritize accordingly.

Total portfolio risk budgets are further complicated by decisions around ex-ante vs. ex-post measurements, sources of risk estimates, and treatment of private markets (e.g., smoothed vs. "economic reality" pricing), for example.

A useful concept for managing the risk budget is marginal contribution to tracking error (MCTE), the amount by which each position increases the total fund's tracking error. Positions with high MCTE are consuming disproportionate shares of the budget. If a position has high MCTE but modest expected risk-adjusted return, it is a candidate for reduction. If it has high MCTE and high expected risk-adjusted return, it may be worth its cost.

This framing whereby each active bet must justify the risk budget it consumes is the operational heart of TPA's competition-for-capital logic. It is also why TPA requires sophisticated risk infrastructure; the team needs real-time or near-real-time visibility into MCTE across the whole portfolio, including illiquid holdings.

Factor Exposure Calculation

TPA funds typically decompose portfolio risk into a small number of systematic factors, most commonly some combination of equity market beta, interest rates, credit spreads, inflation, and (for non-US funds) currency. Every investment in the portfolio is then mapped to these factors, producing a single integrated picture of what the portfolio is exposed to.

The choice of factors matters more than it might appear. A framework that uses only equity beta and rates will produce a different picture of the portfolio than one that also includes inflation, credit spreads, or liquidity risk. A fund that is heavily exposed to inflation-sensitive assets would look quite different if inflation is explicitly modeled as a separate factor. Conversely, adding too many factors risks spurious precision. That is, with enough factors, almost any portfolio may look "diversified," and the framework loses its ability to identify genuine concentrations.

For liquid public assets, factor exposures are typically estimated using regression analysis. The return history of a given asset or fund is regressed against the return history of the factor proxies (e.g., a global equity index for the equity beta factor, a long-duration government bond index for the rates factor, a credit spread index for the credit factor). The regression coefficients (betas) represent the asset's sensitivity to each factor.

The regression is typically run over a historical window of three to five years, though the choice of window involves trade-offs. Longer windows capture more data and reduce noise but may include periods when the asset's risk characteristics were materially different. Shorter windows are more responsive to recent changes in a fund's strategy or market structure but are more sensitive to outlier observations. Some practitioners apply exponential weighting to give more influence to recent observations without discarding older data entirely.

Consider an illustrative example: suppose a core bond fund has a regression that produces an equity beta of 0.15, a rates beta of 0.7, and a credit spread beta of 0.4. This tells the risk system that a 10% decline in global equities is associated with roughly a 1.5% decline in this bond fund's value, independent of its interest rate and credit behavior. The fund's credit spread beta also implies that the fund is less sensitive to credit market stress than interest rate changes.

Private assets present a more significant challenge. They do not have daily return histories, their valuations are reported with a lag, and their "true" factor exposures are partially obscured by smoothing in the valuation process. The two most common approaches are:

- **Public Market Equivalent (PME) Mapping.** The private asset is mapped to a public market proxy that is assumed to share its underlying risk characteristics. A buyout fund might be mapped to a levered small cap equity index; a core infrastructure fund to a combination of long-duration bonds and regulated utility equities; a private credit fund to a leveraged loan index. The factor exposures of the proxy are then used as the exposures of the private asset. This approach represents the assignment of "liquid equivalents" to illiquid holdings.

  The choice of proxy is subjective, particularly for assets with hybrid characteristics. The choice will meaningfully affect the estimated factor exposures. In practice, many institutions develop and document a mapping table that assigns standard proxies to each private market strategy type. This table is reviewed periodically and updated when the strategy mix changes materially. The subjectivity does not disappear, but it becomes explicit and auditable.
- **Fundamental Factor Assignment.** Rather than mapping to a proxy, the investment team directly assigns factor sensitivities based on the underlying economics of the asset. A toll road with inflation-linked revenues and long-duration contracted cash flows might be assigned a high inflation beta, moderate rates sensitivity, and low equity beta.

  This approach is more judgment-intensive than PME mapping, and it is worth acknowledging that both methods involve significant subjectivity. The advantage of fundamental factor assignment is that it may better capture the idiosyncratic economics of unusual assets. The two approaches are not mutually exclusive: some funds use PME mapping as a default and fundamental assignment as an override for assets that clearly do not fit any standard proxy.

Once individual factor exposures are estimated, aggregating to the total portfolio level is straightforward in principle. If private equity represents 20% of the portfolio with an equity beta of 1.3, it contributes 0.26 (20% × 1.3) to the total portfolio's equity beta equivalent. All such contributions are summed across asset classes to produce total factor exposures.

The result, sometimes called "equity equivalent exposure" or "equity beta equivalent", gives the investment team a single number representing the portfolio's aggregate sensitivity to equity markets, regardless of where that exposure sits. The investment team uses this figure, along with analogous aggregates for other factors, to monitor whether the portfolio's total risk profile is consistent with the board's stated risk appetite as reflected in the reference portfolio. If the equity beta equivalent drifts materially above the level implied by the reference portfolio, it signals that the portfolio is taking on more equity-like risk than the board authorized, and the team must consider whether to reduce exposure or seek board guidance.

Illiquidity: Accounting for What Cannot Be Traded

The standard approach for private market assets is to "unsmooth" reported returns by adjusting for the serial correlation that is endemic to the valuation process. Private market prices tend to move slowly and smoothly even when the underlying economic value of the assets is more volatile, a phenomenon well-documented in the real estate and private equity literature.[27] Unsmoothing techniques attempt to reconstruct what the mark-to-market volatility of the asset would look like if it were priced continuously.

The practical effect is to increase the estimated volatility, and therefore the risk budget consumption, of private market assets relative to what their reported returns would suggest. This is the intellectually honest approach: it prevents the portfolio from appearing less risky than it is simply because private assets are valued infrequently.

A simpler but cruder approach is to apply a volatility "gross-up" factor to private market exposures, assuming, for example, that the true volatility of a private equity portfolio is 1.5-2x the volatility observed from smoothed quarterly returns. This is less precise but operationally easier.

For new funds or co-investments with no return history, most institutions fall back on an asset class proxy, using the historical returns of a comparable strategy type as a stand-in until sufficient track record accumulates. Some also use the GP's prior fund history if it is available and deemed representative. Neither approach is fully satisfactory for a genuinely novel investment, and risk systems that rely on historical data will always have a "new investment" problem.

Because private market assets cannot be rebalanced quickly, TPA funds typically manage illiquidity risk through pacing and liquidity forecasting rather than through tactical adjustment. The investment team models the expected trajectory of private market NAVs, distributions, and capital calls over a multi-year horizon, and manages the liquid portion of the portfolio to ensure the fund can meet obligations (benefit payments, capital calls) without forced selling.

The deeper implication is that illiquid assets effectively "lock in" a portion of the risk budget for years at a time. A commitment to a private equity fund made today will consume tracking error for the next ten to fifteen years, largely regardless of what the investment team does. This makes pacing decisions, such as how much to commit each year and across which strategies, a more consequential part of TPA portfolio management than they were under SAA. This is because under SAA, if the private equity allocation drifts above its target, the policy bands themselves create a natural brake on new commitments. Under TPA, there is no such automatic brake; the investment team must actively manage the pace of commitments as part of its risk budget management, making pacing a more explicit and consequential judgment call.

Currency: Hedging Policy and Active Management

Most large TPA funds separate the currency decision into two components: a strategic (or passive) hedging policy and an active currency overlay.

The strategic hedging policy defines a baseline hedge ratio, the percentage of foreign currency exposure that is systematically hedged back to the base currency regardless of market views. A fund might, for example, adopt a policy of hedging 50% of all developed market foreign currency exposure and 0% of emerging market exposure (where hedging costs are higher). This policy is set at the total portfolio level, not sleeve by sleeve.
On top of this baseline, some TPA funds run an active currency overlay, a program of tactical over- and under-hedging based on views about factors such as currency valuations, interest rate differentials, and momentum.

In a TPA framework, currency exposure, including both the unhedged portion of foreign asset holdings and any active currency positions, consumes part of the active risk budget. An unhedged international equity allocation, for example, contributes both equity beta (from the underlying equities) and currency beta (from the unhedged exchange rate exposure) to total fund tracking error.

The cost of hedging, the difference between the forward rate and the spot rate, reflects the interest rate differential between the two currencies (i.e., covered interest parity). This hedging cost is not trivial and must be factored into the expected return calculation for international assets. An international equity allocation that generates 7% in local currency terms but costs 150 basis points to hedge back to dollars has an expected net return of 5.5%, and that cost must be weighed against the benefit of reduced currency volatility. TPA funds that manage currency actively are, in effect, making a judgment about when the cost of hedging is worth paying and when it is not.

A Note on Integration

The four components described above are not independent systems. They interact, hence a single integrated risk platform may prove more helpful than four separate tools. Fully integrated platforms that handle public markets risk, private markets factor mapping, liquidity forecasting, and currency overlay in a single system remain rare. Most institutions assemble this capability from multiple vendors[28] and a custom or third-party aggregation layer. The integration work (of matching data formats, reconciling valuation methodologies, and building consistent reporting) is typically done in-house or with the help of a specialist consultant. No single vendor has yet delivered a solution that institutional investors broadly regard as having fully solved this problem.

A private equity commitment, for example, affects the risk budget (through its equity beta equivalent contribution and its illiquidity-adjusted volatility), the factor exposure picture (through its growth and credit factor loadings), and potentially the currency picture (if the fund invests in non-USD assets). Changes in currency hedging affect the factor exposure picture (by changing the portfolio's sensitivity to foreign equity returns) and consume risk budget. The whole point of TPA is that these interactions are managed at the total portfolio level, which is why the technology and data infrastructure challenge is as significant as it is.

No institution, including the most sophisticated TPA practitioners, has perfectly solved this integration problem. The goal is not perfection but sufficiency: a risk framework that is good enough to catch the major risk exposures, flag the significant budget overruns, and support better decisions.

Appendix: TPA as OCIO

For smaller institutions that lack the internal investment staff to implement TPA directly, an outsourced CIO (OCIO) relationship offers a potential path to some of TPA's benefits without the full infrastructure build. The fit is imperfect, and the limitations are real, but they are worth understanding clearly.

A well-resourced OCIO may be able to fulfill several of TPA's core technical requirements that smaller institutions cannot easily replicate internally. These could include a total-portfolio risk framework spanning public and private markets, active risk budgeting at the total fund level, factor-based exposure monitoring, integrated liquidity forecasting, and dynamic reallocation across asset classes in response to market conditions. For an institution with no dedicated internal investment team, these capabilities may simply be unavailable any other way.

End Notes

[1] See Harry Markowitz, "Portfolio Selection," The Journal of Finance 7, no. 1 (1952): 77–91.

[2] Source: Pensions & Investments, "CalPERS CIO Stephen Gilmore looks to import the total portfolio approach," Arleen Jacobius, November 19, 2024.

[3] Sources: New Zealand Superannuation Fund, "Actual Portfolio," https://nzsuperfund.nz/how-we-invest/actual-portfolio; IFSWF, "New Zealand Superannuation Fund," https://www.ifswf.org/member-profiles/new-zealand-superannuation-fund-0.

[4] Source: CalPERS, Investment Committee Agenda Item 5a, Attachment 1, November 2025.

[5] See Meketa's Risk Budgeting research paper for a more detailed explanation of the active risk budgeting concept and its implementation.

[6] Source: CalPERS, Investment Committee Agenda Item 5a, Attachment 1, November 2025.

[7] Source: GIC, "GIC's Total Portfolio Approach," GIC ThinkSpace, August 22, 2022.

[8] Source: CPP Investments, "Our Investment Strategy," https://www.cppinvestments.com/the-fund/how-we-invest/our-investment-strategy/. CPP has two separate reference portfolios reflecting the different risk profiles of the base and additional CPP accounts. The base CPP reference portfolio is 85% global equity (MSCI World Index) and 15% Canadian nominal government bonds. The additional CPP reference portfolio, which carries greater sensitivity to investment shortfalls given its different funding profile, is 55% global equity and 45% Canadian government bonds.

[9] The Future Fund does not use a reference portfolio in the TPA sense. Its benchmark is set by government mandate as CPI plus 4% to 5% per annum over the long term, with an acceptable but not excessive level of risk. This is a real return target rather than a passive investable benchmark.

[10] Source: GIC Report 2024/25, https://report.gic.com.sg/investment-report.html. GIC's reference portfolio is composed of 65% global equities and 35% global bonds. GIC explicitly states that "the Reference Portfolio is not a performance benchmark for the GIC Portfolio." Rather, it represents the risk level the Government of Singapore is prepared for GIC to take.

[11] The CalPERS Board voted to adopt TPA in November 2025, with implementation beginning July 1, 2026.

[12] Sources: Top1000funds.com, "NZ Super Co-CIOs Chart TPA Vision; Hunt for New Alpha Sources," September 16, 2025; Financial Times, "The hot new investment trend is the 'Total Portfolio Approach'. Does it work?", Toby Nangle, November 17, 2025; author discussion with Stephen Gilmore.

[13] Source: CAIA Association, "Innovation Unleashed," 2024.

[14] CPP Investments, "CPP Investments Net Assets Total $632.3 Billion at 2024 Fiscal Year End," press release, May 22, 2024, https://www.cppinvestments.com/newsroom/cpp-investments-net-assets-total-632-3-billion-at-2024-fiscal-year-end/. As of fiscal year 2024, CPP Investments reported a 10-year annualized net return of 9.2%. However, CPP disclosed that over the trailing five-year and ten-year periods, net value-added relative to the reference portfolios was negative 2.0% and negative 0.3%, respectively.

[15] Source: CPP Investments, "Statement of Investment Objectives, Policies, Return Expectations and Risk Management," April 2, 2025. The Board annually reviews and approves market risk limits for each Investment Portfolio, expressed in equity/debt risk equivalency terms and absolute potential loss (value at risk) terms.

[16] In practice, it is advisable to establish decision rights before adopting TPA: who sets the reference portfolio, who can revise it (and how often), what actions require escalation, and how exceptions are handled. Without that wiring, TPA can become a new label applied to old habits.

[17] Source: Matt Toledo, "Implementing Total Portfolio Approach at an Operational Level," Chief Investment Officer, January 20, 2026.

[18] Source: Campbell R. Harvey, Yan Liu, and Heqing Zhu, "...and the Cross-Section of Expected Returns," Review of Financial Studies 29, no. 1 (2016): 5–68.

[19] Source: Texas Teachers Retirement System, "Special Topic: Total Portfolio Approach," February 2026.

[20] Source: Darcy Song, "Texas Teachers' CIO Questions TPA, DAA Value-Add," Top1000funds.com, February 19, 2026.

[21] Source: Darcy Song, "Texas Teachers' CIO Questions TPA, DAA Value-Add," Top1000funds.com, February 19, 2026.

[22] Source: Richard Grinold, "The fundamental law of active management", Journal of Portfolio Management, Spring 1989.

[23] TPA can be set up to evaluate performance on two levels. The first, and crucial, question is whether the team made effective asset selections compared to the reference portfolio. The second question is how well the (asset class) team executed their strategy, which can be measured against an asset class benchmark.

[24] Source: Willis Towers Watson / Thinking Ahead Institute, "The Total Portfolio Approach," November 2020; CalPERS press release, November 2025.

[25] Source: CAIA Association, "Beyond Comparison — Total Portfolio Approach and Performance Analysis," Chronicles of an Allocator, March 2024.

[26] Source: PitchBook, "CalPERS' Unquantifiable Overhaul," November 25, 2025.

[27] See, for example, Geltner, David M. "Estimating Market Values from Appraised Values without Assuming an Efficient Market." Journal of Real Estate Research 8, no. 3 (1993): 325–346 and Ang, Andrew, Dimitris Papanikolaou, and Mark M. Westerfield. "Portfolio Choice with Illiquid Assets." Management Science 60, no. 11 (November 2014): 2737–2761.

[28] For example, Axioma, MSCI Barra, or BlackRock Aladdin may be used for a public markets risk system, and Burgiss, Preqin, or FactSet may be used as a private markets data provider.

Important Information

Quotations from and references to third-party individuals and institutions reflect those parties' views and do not represent the views or recommendations of Meketa Investment Group.

This report (the "report") has been prepared for the sole benefit of the intended recipient (the "recipient").

Significant events may occur (or have occurred) after the date of this report, and it is not our function or responsibility to update this report. The information contained herein, including any opinions or recommendations, represents our good faith views as of the date of this report and is subject to change at any time. All investments involve risk, and there can be no guarantee that the strategies, tactics, and methods discussed here will be successful.

The information used to prepare this report may have been obtained from investment managers, custodians, and other external sources. Some of this report may have been produced with the assistance of artificial intelligence ("AI") technology. While we have exercised reasonable care in preparing this report, we cannot guarantee the accuracy, adequacy, validity, reliability, availability, or completeness of any information contained herein, whether obtained externally or produced by the AI.

The recipient should be aware that this report may include AI-generated content that may not have considered all risk factors. The recipient is advised to consult with their Meketa advisor or another professional advisor before making any financial decisions or taking any action based on the content of this report. We believe the information to be factual and up-to-date but do not assume any responsibility for errors or omissions in the content produced. Under no circumstances shall we be liable for any special, direct, indirect, consequential, or incidental damages or any damages whatsoever, whether in an action of contract, negligence, or other tort, arising out of or in connection with the use of this content. It is important for the recipient to critically evaluate the information provided.

Certain information contained in this report may constitute "forward-looking statements," which can be identified by the use of terminology such as "may," "will," "should," "expect," "aim," "anticipate," "target," "project," "estimate," "intend," "continue," or "believe," or the negatives thereof or other variations thereon or comparable terminology. Any forward-looking statements, forecasts, projections, valuations, or results in this report are based upon current assumptions. Changes to any assumptions may have a material impact on forward-looking statements, forecasts, projections, valuations, or results. Actual results may therefore be materially different from any forecasts, projections, valuations, or results in this report.

Performance data contained herein represent past performance. Past performance is no guarantee of future results.